

Call for participation

קול קורא להשתתפות

Dear friends and colleagues,

Join us for GSC'26—the annual gathering for graduate students in Systems and Control Theory! The attached program features presentations by 7 Ph.D. and 11 M.Sc. students from various universities across Israel, offering high-level insights and thought-provoking ideas spanning a plethora of topics in Systems and Control. Do not miss the chance to network with peers and immerse yourself in the latest advances in our field!

This year the event will be hosted on Monday, April 27, 2026, by the School of Electrical and Computer Engineering at Tel Aviv University. It is scheduled to take place in the Kitot building, room 011, which is located next to the School of Social Work and Gate 17. The [Google Maps link](#) shows the walking path from the Tel Aviv University train station, as in the map below.

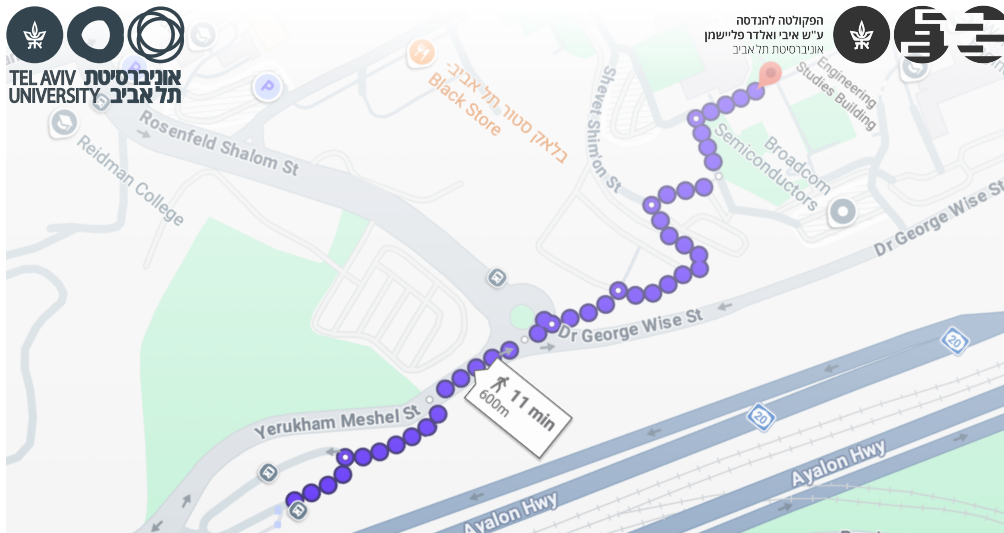
To help us to plan, please register at <https://forms.gle/WBdgnqGGMHziNAKaA> by **April 22, 2026, 8:00pm**.

We gratefully thank Tel Aviv University, and especially the Faculty of Engineering, for their generous comprehensive support and assistance, as well as for partially funding the event.

Looking forward to seeing you at GSC'26,

RAMI KATZ & ANATOLY KHINA
LEONID MIRKIN

GSC'26 Organizers
IAAC President



National Member Organization of [IFAC](#) and [IAIN](#)

Invitation to IAAC workshop

Graduate Students in Systems and Control

GSC'26

to be held in **room 011 Engineering Classroom bld.**, Tel Aviv University
on Monday, April 27, 2026 (Iyar 10, 5786)

Organizers: **Rami Katz & Anatoly Khina** (TAU)

We are grateful to the organizations below, whose support makes holding IAAC events possible

Applied Materials Israel Ltd.

Cielo Inertial Solutions Ltd.

Elbit Systems Ltd.

RAFAEL—Advanced Defense Systems Ltd.

Program

תכנית

08:30–08:55	Gathering
08:55–09:00	Opening remarks
09:00–09:20	Adam Kaminer (ECE@TAU; supervisor: M. Margaliot) <i>A Turnpike Property in an Eigenvalue Optimization Problem</i>
09:20–09:40	Omer Wexler (AE@Technion; supervisor: M. Idan) <i>Pointwise Optimal Control for an Air-Breathing Hypersonic Vehicle with Multiple Scramjet-Imposed Constraints</i>
09:40–10:00	Lidor Erez (IE@BGU; supervisors: A. Taitler & S. Shperberg) <i>From Kinematics to Dynamics: Learning to Refine Hybrid Plans for Physically Feasible Execution</i>
10:00–10:20	Eyal Bar-Shalom (ECE@TAU; supervisor: M. Margaliot) <i>Guardian Maps for Continuous-Time Systems: A Lie-Algebraic Approach</i>
10:20–10:40	Amit Gedj (ME@BGU; supervisor: S. Arogeti & A. Taitler) <i>Disturbance-Aware Data-Driven Optimal Altitude Control of UAVs</i>
10:40–11:00	Coffee / tea break
11:00–11:20	Nadav Barak (ME@Technion; supervisor: C. Grubler) <i>MFDOOM: Model-Free Disturbance Compensation for Data-Enabled Predictive Control</i>
11:20–11:40	Zamir Martinez (AE@Technion; supervisor: D. Zelazo) <i>Formation Control via Rotation Symmetry Constraints</i>
11:40–12:00	Zvi Chapnik (ME@Technion; supervisors: Y. Or & S. Revzen) <i>Sample Efficient Learning of Body-Environment Interaction of an Under-Actuated Locomotion System</i>
12:00–12:20	Rotem Kain (CEE@Technion; supervisor: A. Degani) <i>Error-State Quaternion EKF for Pose and Shape Estimation of a Continuum Robot Using Internal Vision-Based Sensing</i>

12:20–12:40	Kang Tong (AE@Technion; supervisors: C. Grubler & M. S. Chong) <i>Periodic Fixed-Points and Their Algebraic Characteristics in Discrete-Time Lur'e Feedback Systems</i>
12:40–12:45	2025 IAAC Award ceremony
12:45–14:15	Lunch break
14:15–14:35	Ilan Minkin (Math@TAU; supervisor: E. Fridman) <i>Constructive Method for Averaging of Singularly Perturbed Linear Systems with Rapidly Varying Coefficients</i>
14:35–14:55	Zeev Yampolsky (MT@UHaifa; supervisor: I. Klein) <i>Improving Underwater Navigation Through Accurate Doppler Velocity Log Calibration</i>
14:55–15:15	Adir Goldovsky (ECE@TAU; supervisor: G. Weiss) <i>Synchronverters with Adaptive Harmonic Mitigation Capability</i>
15:15–15:35	Zijian Liu (AE@Technion; supervisor: C. Grubler) <i>Verification of Variation Bounding in Discrete-Time Hankel Operators</i>
15:35–15:55	Coffee / tea break
15:55–16:15	Fengyu Yue (AE@Technion; supervisor: D. Zelazo) <i>Transversal Passivity: A Geometric Extension of Passivity Theory</i>
16:15–16:35	Harutjun Magakyan (ECE@Technion; supervisor: N. Shimkin) <i>Anomaly Detection in Filtered Sequences via Conditional Variational Encoding</i>
16:35–16:55	Tal Neshet (ME@Technion; supervisor: Y. Or) <i>Experimental Investigation of Optimal Control for Planar Three-Link Swimmer in Viscous Fluid</i>
16:55–17:15	Solomon Goldgraber Casspi (AE@Technion; supervisor: D. Zelazo) <i>The Geometry of Hidden Modes in Distance-Based Formation Control</i>
17:15–17:20	Closing remarks

09:00–09:20

*A Turnpike Property in an Eigenvalue Optimization Problem***Adam Kaminer** (ECE@TAU; supervisor: M. Margaliot)**Abstract:**

We consider a constrained eigenvalue optimization problem that arises in an important nonlinear dynamical model for mRNA translation in the cell. We prove that the ordered list of optimal parameters admits a turnpike property, namely, it includes three parts with the first and third part relatively short, and the values in the middle part are all approximately equal. Turnpike properties have attracted considerable attention in econometrics and optimal control theory, but to the best of our knowledge this is the first rigorous proof of such a structure in an eigenvalue optimization problem. For more details, see <https://arxiv.org/html/2601.13756v1>

09:20–09:40

*Pointwise Optimal Control for an Air-Breathing Hypersonic Vehicle with Multiple Scramjet-Imposed Constraints***Omer Wexler** (AE@Technion; supervisor: M. Idan)**Abstract:**

Air-breathing hypersonic vehicles (AHVs) in both civilian and military fields are attracting the attention of many industries, research institutes, and countries alike. Such aerial vehicles are highly unstable and nonlinear in nature. Moreover, the efficient operation of scramjet engines that normally power the AHVs imposes strict constraints on the flight conditions of the vehicle. These intricate characteristics, alongside their promising capabilities, incite interest in this type of aircraft.

This work aims to address the tracking control problem of a generic AHV with multiple state constraints related to its propulsion system. It is assumed that the vehicle can maneuver using an elevator and thrust control attained by varying the stoichiometrically normalized fuel-to-air ratio. The dynamics of this aircraft are characterized via a complicated “truth” model, which is not suitable for control design. A curve-fitted model (CFM) is derived from the truth model to capture its essence. Finally, a control-oriented model (COM) is obtained from the CFM by assuming constant flight altitude and setting the weak elevator couplings to zero.

A pointwise-optimal controller is developed in this study for the COM, accounting for both tracking performance requirements and engine-imposed state-dependent constraints. Those two requirements are expressed using appropriate Lyapunov functions: a quadratic control Lyapunov function (CLF) and a barrier Lyapunov function (BLF), respectively. The controller also addresses modelling errors (model uncertainty) arising from the difference between the CFM and COM. Its design can be greatly simplified by adding a canard control surface to the AHV. A configuration without a canard results in a higher relative degree of the commanded variables and BLF. The current study addressed this case. The high relative degree and structure of the

model require the use of more advanced controller design tools compared to the case with a canard addressed previously. Exponential BLFs (EBLF) are utilized to comply with generic state-dependent constraints. In addition, the EBLF control logic is modified to provide robustness to modelling errors. The design is validated through numerical simulation with a realistic maneuver of the AHV. The accurate tracking performance of the proposed controller while complying with the state-dependent engine-imposed constraints also demonstrates its robustness to modelling errors and disturbances.

09:40–10:00

*From Kinematics to Dynamics: Learning to Refine Hybrid Plans for Physically Feasible Execution***Lidor Erez** (IE@BGU; supervisors: A. Taitler & S. Shperberg)**Abstract:**

In many robotic tasks, agents must traverse a sequence of spatial regions to complete a mission. Such problems are inherently mixed discrete-continuous: a high-level action sequence and a physically feasible continuous trajectory. The resulting trajectory and action sequence must also satisfy problem constraints such as deadlines, time windows, and velocity or acceleration limits. While hybrid temporal planners attempt to address this challenge, they typically model motion using linear (first-order) dynamics, which cannot guarantee that the resulting plan respects the robot’s true second-order physical constraints. Consequently, even when the high-level action sequence is fixed, producing a dynamically feasible trajectory becomes a bi-level optimization problem. We address this problem via reinforcement learning in continuous space. We define a Markov Decision Process that explicitly incorporates analytical second-order constraints and use it to refine first-order plans generated by a hybrid planner. Our results show that this approach can reliably recover physical feasibility and effectively bridge the gap between a planner’s initial first-order trajectory and the dynamics required for real execution.

10:00–10:20

*Guardian Maps for Continuous-Time Systems: A Lie-Algebraic Approach***Eyal Bar-Shalom** (ECE@TAU; supervisor: M. Margaliot)**Abstract:**

Guardian maps are scalar maps that vanish when a matrix or polynomial is on the verge of stability, and thus provide a useful tool for stability analysis. Several guardian maps have been proposed in the literature for Hurwitz stability based on the Kronecker sum, the second lower Schläflian matrix, and the bialternate sum. It is natural to seek a unifying principle underlying all these maps. We introduce the novel Lie-algebraic notion of a guardian representation, and present the Kronecker sum and the second additive compound as representative examples of this framework. We demonstrate an application of our theoretical results for deriving guardian maps for discrete-time stability.

10:20–10:40

Disturbance-Aware Data-Driven Optimal Altitude Control of UAVs

Amit Gedj (ME@BGU; supervisor: S. Arogeti & A. Taitler)

Abstract:

Adaptive dynamic programming (ADP) and policy iteration (PI) algorithms are powerful tools for data-driven optimal control. When using the linear quadratic regulator (LQR) in the framework of ADP, most existing approaches assume disturbance-free system dynamics. Data-driven design eliminates the need for a dynamical model by utilizing data from the system state and input. If the system is disturbed by an unknown disturbance, it means the system input is not entirely known, which makes standard ADP and PI approaches impractical. In this study, a novel approach for the PI-based ADP framework is proposed to explicitly handle constant unknown disturbances by integrating frequency-domain filtering with the classical ADP approach. High-pass filtering of input-output trajectories suppresses the effects of disturbances on the design, allowing the usage of the already established PI techniques, while a low-pass filter extracts the required steady-state, feed-forward term. The resulting control law preserves the LQR structure but improves steady-state accuracy and robustness. The process is experimentally validated using a quadcopter altitude control design with unknown mass and motor dynamics. The unknown gravitational force is assumed to be a constant unknown disturbance input. The results show that the proposed approach eliminates steady-state altitude errors compared to baseline controller gains. These findings demonstrate that disturbance-aware ADP provides a practical and effective framework for robust data-driven control in systems with unknown dynamics.

11:00–11:20

MFDOOM: Model-Free Disturbance Compensation for Data-Enabled Predictive Control

Nadav Barak (ME@Technion; supervisor: C. Grubler)

Abstract:

Data-Enabled Predictive Control (DeePC) enables control of unknown systems directly from data, eliminating the need for explicit modeling. Yet in practice, its performance hinges on a critical assumption: that past data remains representative of future behavior. When disturbances or unmodeled dynamics are present, this assumption breaks down, leading to degraded predictions and compromised control performance.

In this talk, we challenge this limitation by asking: can a controller learn disturbance behavior directly from its own prediction errors? We introduce Model-Free Disturbance Observer with Online Modification (MFDOOM), a framework that augments DeePC by continuously learning from past prediction mismatches. MFDOOM constructs and updates a Hankel matrix of prediction errors, using it to correct future predictions in real time. For disturbances that can be represented as the output of an autonomous linear time-invariant system, the method enables prediction of future errors without relying on persistent excitation.

This perspective turns prediction errors into a resource: a structured signal that implicitly encodes disturbance dynamics. We compare MFDOOM to recent online and time-varying DeePC

methods, highlighting its ability to adapt with reduced data requirements while improving prediction accuracy and closed-loop performance.

The proposed approach suggests a shift in how disturbance rejection can be achieved in data-driven control: not by explicitly modeling disturbances, but by exploiting the structure of the errors they leave behind.

11:20–11:40

Formation Control via Rotation Symmetry Constraints

Zamir Martinez (AE@Technion; supervisor: D. Zelazo)

Abstract:

This work introduces a distributed formation control strategy for multi-agent systems based solely on rotation symmetry constraints. We propose a potential function that enforces inter-agent **rotational** symmetries, whose gradient defines a control law that drives the agents toward a desired planar symmetric configuration. We show that only $n - 1$ edges—the minimal connectivity requirement—are sufficient to implement the proposed strategy, where n is the number of agents. We further augment the design to address the **maneuvering problem**, enabling the formation to undergo coordinated translations, rotations, and scaling along a predefined virtual trajectory. Simulation examples are provided to validate the effectiveness of the proposed method.

11:40–12:00

Sample Efficient Learning of Body-Environment Interaction of an Under-Actuated Locomotion System

Zvi Chapnik (ME@Technion; supervisors: Y. Or & S. Revzen)

Abstract:

Geometric mechanics provides a principled framework for understanding how biological and robotic systems generate motion through controlled cyclic changes in body shape. In friction-dominated regimes, this interaction can often be described by a reduced representation in which body velocity depends directly on the controlled shape and shape velocity, commonly referred to as the motility map. This formulation enables analysis of locomotion in terms of cyclic shape changes, where net displacement depends primarily on the geometry of the gait rather than its execution rate. However, learning such models from experimental data remains challenging, particularly in systems with underactuation and complex body-environment interactions.

In this work, we investigate data-driven identification of motility maps for a physical robotic system designed to emphasize these challenges. The platform is a multi-link swimmer with two actuated joints and four passive flippers operating in a granular substrate, introducing both shape-underactuation and highly nonlinear interactions that are difficult to model analytically.

We compare modeling approaches that differ in data expressiveness and in how they account for passive dynamics, ranging from linear, phase-dependent models to nonlinear Gaussian mixture regression models, as well as extensions that incorporate reduced representations of passive

shape dynamics. Performance of the different models is evaluated by predicting body velocity across multiple gaits and actuation speeds, enabling assessment of both interpolation and generalization.

The results reveal a trade-off between model simplicity and accuracy: simpler models perform well with limited data, while more expressive models improve with additional data and achieve lower prediction error. All methods outperform the baseline phase-dependent model, with the greatest improvements achieved by nonlinear models that incorporate passive dynamics when sufficient training data are available. These findings suggest that structure-preserving, data-driven models can support accurate prediction and provide a foundation for feedback control and adaptation in underactuated systems with complex body-environment interactions.

12:00–12:20

Error-State Quaternion EKF for Pose and Shape Estimation of a Continuum Robot Using Internal Vision-Based Sensing

Rotem Kain (CEE@Technion; supervisor: A. Degani)

Abstract:

Continuum robots are well suited for operation in confined environments, but their deformable structure makes reliable pose and shape estimation challenging without external sensing. This work presents an internal vision-based estimation framework using only onboard measurements. An internal actuator equipped with monocular cameras provides visual motion cues, complemented by encoder and gyroscope data.

The actuator motion inside the robot is modeled using a unicycle kinematic model, consistent with the available forward velocity and angular-rate measurements, and used as the prediction model within an Extended Kalman Filter (EKF). A 2D formulation is first used to validate visual odometry and multi-rate sensor fusion, and is then extended to full 3D motion.

Using Euler angles leads to instability near singular configurations due to gimbal lock, while a nominal quaternion EKF removes the singularity but can still become unstable under aggressive rotations due to inconsistent additive updates. To address this, we formulate an Error-State Quaternion EKF (ES-Q-EKF), where the orientation is represented with a quaternion and only small rotation errors are estimated and applied multiplicatively.

Simulation results show that the ES-Q-EKF eliminates gimbal-lock failure and significantly improves stability and accuracy compared to both Euler-angle and nominal quaternion formulations, enabling reliable pose and shape estimation using only onboard sensing.

12:20–12:40

Periodic Fixed-Points and Their Algebraic Characteristics in Discrete-Time Lur'e Feedback Systems

Kang Tong (AE@Technion; supervisors: C. Grubler & M. S. Chong)

Abstract:

We study the problem of determining nontrivial, i.e., non-constant, periodic fixed points in

discrete-time Lur'e feedback systems. Using the eigenvalues of the circulant matrix constructed from the impulse response of the linear subsystem, we introduce an algebraic framework that allows us to determine the existence of such fixed-points. Concretely, we consider the intersection of the graph of the static nonlinearity and the sector bound, whose slopes are generated by the largest and smallest modulus of the inverses of the nonzero eigenvalues. For static nonlinear functions that pass through the origin, we show that a necessary condition for the system admitting nontrivial P -periodic fixed-points is that the intersection contains at least one point other than the origin. Our characterization further allows us to derive bounds on the amplitudes of any admissible periodic fixed-point.

14:15–14:35

Constructive Method for Averaging of Singularly Perturbed Linear Systems with Rapidly Varying Coefficients

Ilan Minkin (Math@TAU; supervisor: E. Fridman)

Abstract:

In this paper we study the averaging-based stability of a class of linear periodic singularly perturbed systems (SPSs) with fast varying coefficients. Under assumption that the fast system and the averaged slow system are stable, we present constructive conditions that guarantee the stability of the SPS for small enough values of singular perturbation parameter ε and averaging parameter μ and provide numerical bounds on these parameters. We apply a recently introduced delay-free transformation for averaging to the slow system and apply an appropriate ε -dependent Lyapunov function to the transformed system leading to efficient LMI conditions for finding the upper bounds on the small parameters that guarantee the stability. We further apply results to stabilization of affine SPS by fast periodic switching, where practical stability is achieved if μ is smaller than ε . We extend the results to SPS with delay in the slow variable, where we employ appropriate Lyapunov-Krasovskii functionals. Numerical examples illustrate the efficiency of our results.

14:35–14:55

Improving Underwater Navigation Through Accurate Doppler Velocity Log Calibration

Zeev Yampolsky (MT@UHaifa; supervisor: I. Klein)

Abstract:

Underwater navigation is inherently challenging, as radio signals cannot penetrate the water body, rendering radio-based aiding sources, such as global navigation satellite systems (GNSS) position updates, unavailable. Consequently, underwater navigation solutions rely on the fusion of inertial sensors with an acoustic aiding sensor, commonly a Doppler velocity log (DVL), which provides velocity measurement. This fusion is typically implemented using an extended Kalman filter (EKF), where DVL velocity is used as an external update to the navigation filter. However, to achieve and maintain accurate performance, DVL calibration must be performed. In underwater marine robotics, autonomous underwater vehicles (AUVs) and other platforms typically conduct this calibration at the surface, where GNSS signals provide an accurate ve-

locity reference. During calibration, the platform must execute specific maneuvers to increase observability and enable accurate estimation of DVL error terms. Estimation filters, including linear and extended Kalman filters, are commonly employed for this purpose. Overall, this leads to a complex and time-consuming calibration procedure.

To reduce the complexity of DVL calibration, we propose two main approaches: a model-based EKF and a deep-learning pipeline, both requiring only a simple constant-velocity trajectory during calibration. We introduce an end-to-end data-driven framework to estimate a comprehensive DVL error model for improved calibration accuracy. We further propose a zero-velocity update with a linear Kalman filter utilizing only four states. We demonstrate increased observability and an improvement of 15% in accuracy compared to the baseline methods.

Additionally, to further investigate the DVL error models used in calibration, we propose DCNet, a data-driven framework for DVL calibration. DCNet is a multi-head architecture consisting of a two-dimensional convolutional neural network (2DCNN) and a one-dimensional convolutional neural network (1DCNN), designed to extract hidden dependencies between the reference and DVL velocity measurements for precise error-term estimation. By employing DCNet, we demonstrate a reduction in calibration time of up to 80%, while achieving an average accuracy improvement of 70%, all under constant-velocity dynamics. These improvements are enabled by estimating a comprehensive six error term model. Such DCNet improvements are achieved for lower-cost, lower-accuracy DVLs.

Both of our proposed approaches were validated on real-world data collected in two sea experiments carried out by the University of Haifa's AUV. This demonstrates the applicability, effectiveness, and robustness of our proposed approaches to achieve accurate DVL calibration even with low-cost DVLs.

14:55–15:15

Synchronverters with Adaptive Harmonic Mitigation Capability

Adir Goldovsky (ECE@TAU; supervisor: G. Weiss)

Abstract:

Harmonic distortion is a growing concern in modern electrical grids due to the widespread use of nonlinear loads and power-electronic converters. Our work presents a harmonic mitigation technique suitable for virtual synchronous machines (VSMs). This technique implements a virtual RLC passive filter for each harmonic to be suppressed inside the VSM algorithm. The resonant frequencies of the harmonic suppression mechanism are adjusted adaptively according to the grid frequency, which is well estimated by the VSM internal frequency. The same method can deal also with negative sequence voltages. Simulation results confirm that our technique can significantly reduce voltage harmonics at the inverter terminals, without affecting stability under diverse operating conditions.

15:15–15:35

Verification of Variation Bounding in Discrete-Time Hankel Operators

Zijian Liu (AE@Technion; supervisor: C. Grubler)

Abstract:

We investigate the k -variation bounding property of the discrete-time Hankel operator, i.e., its invariance under the set of signals with a variation (sign changes) of at most k . It is shown that this property is equivalent to the external positivity of k related linear discrete-time systems, which makes it tractable via numerical and analytic certificates. Our approach extends and unifies earlier investigations on k -variation bounding observability operators, as well as k -variation diminishing Hankel and Toeplitz operators.

15:55–16:15

Transversal Passivity: A Geometric Extension of Passivity Theory

Fengyu Yue (AE@Technion; supervisor: D. Zelazo)

Abstract:

Passivity is a standard tool in nonlinear control analysis and synthesis, particularly for networked systems. Classical state-space passivity notions are typically formulated relative either to the origin, as in standard passivity, or to a specific input-output pair, as in shifted and equilibrium-independent passivity. As a result, they do not directly capture manifold-target control problems such as reduced-attitude coordination and synchronization. In such problems, manifolds represent the desired behavior, with motion along the manifold is permitted, whereas transverse deviations are penalized. Montenbruck *et al.* (TAC, vol. 62, no. 10, pp. 5170–5184, 2017) proposed an insightful input-output framework for submanifold stabilization based on passive relations, but the framework operated at the signal-space level and failed to provide a localized state-space notion of passivity relative to prescribed input-output manifolds.

To fill this gap, we introduce *transversal passivity*, a local state-space notion of passivity defined with respect to input and output manifolds. This notion penalizes only the transverse deviations from prescribed target manifolds. Concretely, we define manifold-error maps based on tubular retractions and incorporate them into local dissipation inequalities. We then establish both necessary conditions and sufficient conditions for a system to exhibit transversal passivity. The resulting framework naturally recovers standard and shifted passivity as special cases. Moreover, by selecting suitable manifolds and storage functions, transversal passivity can characterize systems that inherently violate standard passivity requirements.

We further discuss how transversal passivity extends to feedback interconnections, presenting general conditions that ensure convergence to desired target behaviors. The conditions naturally facilitate a symmetry-based approach to control design. Finally, we demonstrate the efficacy of the framework in the nonlinear synchronization of multi-agent systems over directed graphs.

16:15–16:35

Anomaly Detection in Filtered Sequences via Conditional Variational Encoding

Harutjun Magakyan (ECE@Technion; supervisor: N. Shimkin)

Abstract:

Traditional Time Series Anomaly Detection (TSAD) focuses on identifying irregular temporal

patterns within individual data streams. However, detecting failures in filtering algorithms requires a fundamentally different problem formulation. Here, the anomaly detection task is to identify mapping errors between an algorithm’s inputs and outputs—often with no access to the filter’s internal parameters, structural assumptions, or noise profiles. Because such systems produce strongly correlated, high-dimensional outputs driven by latent dynamics, an effective detector must implicitly recover aspects of the underlying system behavior rather than rely solely on temporal irregularities.

Existing approaches used for this input–output anomaly detection setting are limited. Forecasting-based AD assumes autoregressive temporal structure rather than conditional mapping. Autoencoder-style reconstruction methods struggle with stochastic or multimodal mappings. Recurrent models require long state windows and scale poorly in memory. Diffusion-based conditional generators, while expressive, are computationally prohibitive for real-time monitoring. Among generative models, Conditional VAEs (CVAEs) have emerged as a practical and widely adopted baseline for conditional anomaly detection due to their efficiency, stable training, and ability to model multimodal conditional densities. Yet, CVAEs exhibit two structural weaknesses: (i) because the decoder observes the target output during training, CVAEs can inadvertently reconstruct anomalous outputs too accurately, weakening anomaly scores; and (ii) CVAEs often fail to enforce globally coherent latent geometry, limiting interpretability and impairing recovery of underlying physical structure.

To address these limitations, we introduce the Conditional Variational Encoder (CVE), a two-stage variational architecture built on self/cross-attention mechanisms. CVE is a modification of the CVAE that learns a structured latent representation that mediates the entire mapping $X \rightarrow Z \rightarrow Y$ strictly through the latent variable Z . To train the CVE, we propose a novel objective functional grounded in the Conditional Variational Information Bottleneck (CVIB).

Using the Law of Total Covariance, we analytically show that forcing all information flow through Z separates the learnable system dynamics from exogenous noise, enabling conditional modeling without requiring explicit knowledge of filter parameters. For scalability, CVE tokenizes local temporal windows as spatial features, reducing the quadratic memory cost of attention. A teacher–student training scheme first imposes a strongly regularized latent geometry, then trains a student model to map exclusively from input context to the latent space, preventing target-dependent over-reconstruction during inference.

Beyond detection performance, interpretability is a central objective. We show that CVE’s latent variables correlate with intrinsic simulation parameters, revealing emergent system identification and making the method suitable for safety-critical domains. Experiments on non-linear dynamical systems demonstrate that CVE captures highly correlated input–output mappings, provides interpretable latent structure, and delivers sensitive and computationally efficient mapping-error detection compared to recurrent and diffusion-based baselines.

16:35–16:55

Experimental Investigation of Optimal Control for Planar Three-Link Swimmer in Viscous Fluid

Tal Neshet (ME@Technion; supervisor: Y. Or)

Abstract:

Locomotion of micrometer-sized swimmers is governed by low Reynolds number hydrodynamics, demanding non-reciprocal periodic shape changes for propulsion. While Purcell’s planar three-link swimmer is a canonical theoretical model, validating its analytically derived optimal gaits experimentally remains challenging due to unmodeled real-world fluidic and mechanical complexities. To bridge this gap, we developed an improved design of the untethered, macro-scale robotic surrogate, operating in highly viscous silicone oil to emulate the Stokes regime. This experimental design was constructed as a physical realization of the theoretical model established in previous work of Oren Wiesel and Noam Berkovich Lahav, specifically the central rigid sphere he introduced to represent the added drag of the robot’s central flotation block. This platform allows for precise control over the time dependent shape variables (joint values of the robot) that serve as the primary control inputs for the governing dynamic equations. To reconcile the real robot’s physics, specifically the added drag of its central flotation unit, we utilize a modified theoretical formulation that incorporates a central spherical drag element. In our experiments, the robot was placed in a circular pool of silicone oil (large enough to neglect wall interference). Using a camera system, we tracked the robot’s motion to validate the joint inputs and displacement, allowing comparison between analytical results and real-world data. By tracking our robot, we provide a physical demonstration of these large amplitude gaits. Our results provide an experimental validation of the “Geometric Analysis” framework and optimal control trajectories previously developed by Berkovich Lahav. Our results reveal that the relationship between stroke amplitude and phase is counter-intuitive; in some cases, shifting these variables even causes a reversal in the direction of movement.

16:55–17:15

The Geometry of Hidden Modes in Distance-Based Formation Control

Solomon Goldgraber Casspi (AE@Technion; supervisor: D. Zelazo)

Abstract:

This presentation provides a geometric input-output analysis of hidden modes in distance-based formation control. While the stability and convergence properties of distributed gradient-descent laws are well-understood, analyzing how localized disturbances propagate through a formation requires a formal input-output modeling framework. By studying the linearized system dynamics governed by the configuration-dependent stiffness matrix, we formally characterize the formation’s structural limitations.

We introduce a unified geometric framework to decompose the uncontrollable subspace. Specifically, we demonstrate that uncontrollable rigid-body modes manifest as pure global rotations centered exactly at the actuated node, defining the hidden global rotational subspace. To capture internal deformations that are locally invisible to the actuator, we define the local rotational subspace. Finally, we explore the dynamic consequences of this geometric structure, proving a fundamental “shape recovery dichotomy”. We show that the formation’s ability to reject a localized disturbance and recover its target shape is governed entirely by the geometric alignment of the disturbance with the local component of the hidden rotational mode.