האיגוד הישראלי
לסטטיסטיקה
ומדע הנתונים

الرابطة
الإسرائيلية
للإحصاء وعلم البيانات

The Israel Statistics
and Data Science
Association

הכנס השנתי
2025

ABSTRACTS BOOK

לתוכנית הכנס

חברים יקרים,

ברוכים הבאים לכנס השנתי של האיגוד הישראלי לסטטיסטיקה ומדע הנתונים.

השנה, הכנס חוגג 50 שנה לייסוד האיגוד. חגיגה זו מעומעמת עקב האבל על כל ההרוגים מאז השבעה באוקטובר והדאגה לחטופים שעוד לא שוחררו ולפצועים הזקוקים לרפואת הגוף והנפש.

הכנס, בהשתתפות צוותים רב-מקצועיים בארגונים בישראל, מהווה את הבמה המרכזית בארץ להתעדכנות ולהצגת עבודות איכות ומחקרים. כמדי שנה, הכנס דן בסוגיות המרכזיות של סטטיסטיקה ומדע הנתונים, הפעם בדגש על תיאוריה של בינה מלאכותית וסטטיסטיקה וכן על יישומיהן במגוון שטחים. השתדלנו לגוון, כך שכל המשתתפים ימצאו נושאים בעלי תוכן מסקרן.

היריעה רחבה ויש עניין בכל אחת מההרצאות. מכיוון שקשה לרקוד ביותר מחתונה אחת בו-זמנית, כדי להקל על הבחירה ביניהן, מוגשים לכם בזה תקצירים של ההרצאות שעתידות להינתן בכנס.

אנחנו מאחלים לכם כנס מוצלח. שמחים לראות אתכם כאן איתנו,

פרופ' משה פולק

נשיא האיגוד הישראלי לסטטיסטיקה ומדע הנתונים

Model-Based Deep Learning for Sensing and Imaging: Efficient and Interpretable AI

Yonina Eldar

Deep neural networks have achieved unprecedented performance gains across numerous real-world signal and image processing tasks.

However, their widespread adoption and practical deployment are often limited by their black-box nature – characterized by a lack of interpretability and a reliance on large training datasets.

In contrast, traditional approaches in signal processing, sensing, and communications have long leveraged classical statistical modeling techniques, which incorporate mathematical formulations based on underlying physical principles, prior knowledge, and domain expertise.

While these models offer valuable insights, they can be overly simplistic and sensitive to inaccuracies, leading to suboptimal performance in complex or dynamic real-world scenarios.

This talk explores various approaches to model-based learning which merge parametric models with optimization tools and classical algorithms to create efficient, interpretable deep networks that require significantly smaller training datasets. We demonstrate the advantages of this approach through applications in image deblurring, image separation, super-resolution for ultrasound and microscopy, radar for clinical diagnostics, efficient communication systems, low-power sensing devices, and more. Additionally, we present theoretical results that establish the performance advantages of model-based deep networks over purely data-driven black-box methods.

Short bio: Yonina Eldar is a Professor in the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel where she heads the center for Biomedical Engineering and Signal Processing

and holds the Dorothy and Patrick Gorman Professorial Chair. She is also a Visiting Professor at MIT and Princeton, a Visiting Scientist at the Broad Institute, and an Adjunct Professor at Duke University and was a Visiting Professor at Stanford. She is a member of the Israel Academy of Sciences and Humanities, an IEEE Fellow and a EURASIP Fellow. She received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering from Tel-Aviv University, and the Ph.D. degree in electrical engineering and computer science from MIT. She has received many awards for excellence in research and teaching, including the Israel Prize (2025), Landau Prize (2024), IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014) and the IEEE Kiyo Tomiyasu Award (2016). She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), and the Award for Women with Distinguished Contributions. She was selected as one of the 50 most influential women in Israel, and was a member of the Israel Committee for Higher Education.

She is the Editor in Chief of Foundations and Trends in Signal Processing, a member of several IEEE Technical Committees and Award Committees, and heads the Committee for Promoting Gender Fairness in Higher Education Institutions in Israel

פרופ' אלי קונן

מנהל אגף דימות בשיבא
פרופ' חבר החוג לדימות, אוניברסיטת תל אביב

עולם הדימות הרפואי הינו אחד מתחומי הרפואה הראשונים שעולם הבינה המלאכותית AI נכנס אליו במטרה לעזור (ויש האומרים להחליף) לרדיולוגים המפענחים.

חתן פרס נובל לפיזיקה לשנת 2024 Geoffrey Hinton טען בשנת 2016 ש"תוך 5 שנים לא יצטרכו רדיולוגים בעולם" בזכות כניסת ה-AI לתחום... האם כך?

בהרצאתי אסקור לפניכם את ניסיוני האישי ב-20 השנה האחרונות כמנהל אגף דימות בשיבא, עם דגש על הניסיון המחקרי האקדמי והמסחרי שלנו בתחום עיבוד ממוחשב של תמונה רפואית, נסקור עד כמה התקדם הנושא בעולם כולו, מה כבר מוטמע כיום בעבודה היומיומית בבתי בחולים בארץ ובעולם, וגם מה האכזבות שאנו חווים בתחום.

ארחיב גם על הפוטנציאל המתרחב של שימוש ב-Generative AI במקצועות רפואה אחרים.

ביוגרפיה: פרופ' אלי קונן היה חלוץ בישראל בביצוע בדיקות MRI ו-CT של הלב החל מ-2003. שימש מורה ומדריך לדור שלם של רדיולוגים שהם מומחים מובילים כיום ב-CT וב-MRI של הלב.

במקביל שותף פעיל וחלוץ ב-20 השנים האחרונות במחקר המתמקד בעיבוד תמונה ממוחשב, כולל הקמה של מספר מעבדות מחקר בתחום זה בשיבא, ושת"פ מחקרי עם חוקרים באוניברסיטאות שונות, בעיקר - אוניברסיטת תל אביב.

לינק לפרסומים ב-PUBMED:
מחבר ושותף למעל כ-170 פרסומים בספרות המקצועית ברדיולוגיה, קרדיולוגיה ומחקר AI: https://pubmed.ncbi.nlm.nih.gov/?term=konen+e

I A: Data Science in Transition: From Theoretical Models to Interdisciplinary Applications

Chair: Rotem Dror, *U. of Haifa*

Making Text Generation with LLMs Faster

Michael (Mike) Erlihson

Autoregressive LLM decoding is inherently slow, as each token is generated sequentially, creating a fundamental bottleneck in both latency and resource efficiency. This limitation becomes particularly problematic in real-time applications, large-scale deployments, and scenarios requiring long-context inference.

In this talk, I explore cutting-edge techniques aimed at breaking the sequential barrier and improving decoding speed without compromising output quality. Methods such as speculative decoding leverage lightweight draft models to accelerate generation while maintaining accuracy. Parallel fixed-point methods rethink the iterative token update process, reducing dependencies across steps. Meanwhile, multi-token prediction enables models to generate multiple tokens per step, significantly boosting throughput.

These techniques collectively represent a paradigm shift in LLM inference, pushing beyond the traditional autoregressive constraints. The talk distills insights from about 15 research papers, offering both theoretical perspectives and practical implications for deploying faster, more efficient large language models.

Short bio: Michael (Mike) Erlihson is an Israeli AI expert with a Ph.D. in Mathematics. He is currently the Head of AI at a stealth-mode company and hosts the Data Science Decoded and ExplAInable podcasts. Dr. Erlihson has a rich background in deep learning and data science. He has reviewed over

430 deep learning papers and hosted more than 50 recorded podcasts in the field, building a substantial following of about 60,000 on LinkedIn. He is also the founder of the DeepNightLearners community project, aiming to make deep learning scientific papers more accessible.

During his professional career, Dr. Erlihson has taken on key roles in AI research and leadership, applying his expertise to cutting-edge problems in cybersecurity, deep learning, and applied AI. His work focuses on building AI-driven security solutions, bridging theoretical AI advancements with real-world applications. He is also an advocate for making complex AI topics more digestible, frequently sharing insights and mentoring professionals in the field.

Beyond his professional endeavors, Dr. Erlihson is dedicated to education and knowledge-sharing in the AI community, making complex topics accessible through his various content platforms.

How are multidisciplinary research projects born?

Prof. Mor Peleg

Dept. of Information Systems, University of Haifa

The Data Science (DS) Research Center (DSRC) of the University of Haifa aims to promote multidisciplinary data science research for the benefit of society and the environment. Since its inception six years ago, the center grew to include 64 faculty members, 35 postdocs, and 165 graduate students. To address its mission, the DSRC has been providing incentives and increasing the opportunities for collaboration among core-DS scientists and DS scientists from other research domains who formulate needs that can be addressed by DS methods. The needs drive development of new DS methods, and the researchers from other domains and their students learn how to apply DS methods and tools.

In this talk, I will review our mechanisms for promoting multidisciplinary research and educating non-core DS researchers and research students. I will provide examples of successful and exciting research projects from our center, focusing on real-world needs and data, which led to publications, funded grants, and translation into practice.

Short bio:
Mor Peleg holds a BSc and MSc degrees in biology and PhD (1999) in Information Systems, from the Technion and postdoctoral studies (2003) at Stanford's Medical Informatics program. She is Full Professor at the Department of Information Systems, University of Haifa, which she joined in 2003, Founding Director of the University of Haifa's Data Science Center, founder of the university's Data Science BSc program, and former Chair of the Department of Information Systems.

She is the Editor-in-Chief of the Journal of Biomedical Informatics and is Fellow of the American College of Medical Informatics and of the International Academy of Health Sciences Informatics.

Prof. Peleg is internationally renowned in the area of clinical decision support, with a focus on clinical-guideline based decision support and

using ontologies to integrate patient data with clinical knowledge. She led the large-scale European project MobiGuide, providing sensor-monitoring and evidence-based personalized decision-support to patients any time everywhere. Her current research exploits AI and Big Data for personalizing mental health care for cancer patients (https://capable-project.eu/). Another line of research uses LLMs for extending biomedical ontologies for data integration.

Web page: https://mpeleg.hevra.haifa.ac.il/
LinkedIn: https://www.linkedin.com/in/mor-peleg-1593452/

Learning Heuristics for Combinatorial Optimization Problems with Value-based Reinforcement Learning: An Analysis

Orit Davidovich

Applied Math Research Scientist
IBM Research ISRL

Beginning in the 1990's, considerable empirical work has been carried out to train statistical models, such as neural networks (NNs), as learned heuristics for Combinatorial Optimization (CO) problems. When successful, such an approach eliminates the need for experts to design heuristics per problem type. Due to their structure, many hard CO problems are amenable to treatment through reinforcement learning (RL), and, indeed, we find a wealth of literature training NNs of various architectures using value-based, policy gradient, or actor-critic methods. Results are promising, both in terms of empirical optimality gaps compared to known heuristics as well as in terms of inference runtimes compared to classical solvers. Nevertheless, there has been a paucity of theoretical work justifying the use of RL-based heuristics for CO problems either in terms of convergence or in terms of performance. To this end, we propose a novel analysis of RL value-based methods applied to CO problems, which illuminates the success (and limitations) of the celebrated Deep Q-Network (DQN) algorithm in this problem context. Instead of viewing DQN through the lens of Approximate Policy Iteration (API) or Approximate Value Iteration (AVI), we utilize the perspective of compact approximation schemes, which offers us the great advantage of isolating out the approximation error to focus on the estimation error. We then find that sampling, used in DQN to approximate per-iteration loss, is one of the limiting factors for the success of this method. Indeed, at each iterate we solve a sample average approximation (SAA) problem that notoriously requires large sample sizes, which may approach state space scale, to obtain close to optimal solutions.

Address selection model- From Binary Classification to Learning To Rank

Eldor Zang & Yifat Shani Abuhatzera

CBS

The place of usual residence (address) is a key demographic variable. As Israel prepares for its upcoming administrative census, accurate and complete data sources are essential. In this talk, we present an ML-based methodology that integrates multiple datasets to infer the correct address for each individual. We approach the task from two complementary angles: (1) as a binary classification problem, where we determine whether a candidate address is correct, and (2) as a learning-to-rank (LTR) problem, where we rank all individual addresses to identify the correct one. This approach allows us to optimize the model based on ranking-specific metrics such as Normalized Discounted Cumulative Gain (NDCG), resulting in higher accuracy.

מודל בחירת כתובת - מסיווג בינארי לבעיית דירוג

מקום המגורים הקבוע (כתובת) הוא משתנה ליבה דמוגרפי. לקראת מפקד האוכלוסין המנהלי הקרוב, יש צורך במקורות מידע מלאים ומדוייקים. בהרצאה זו, נציג מתודולוגיה מבוססת למידת מכונה (ML) המשלבת מספר מקורות מידע שונים על מנת לקבוע מהי הכתובת הנכונה עבור כל אדם. נציג את הבעיה בעזרת שני כיוונים משלימים: (1) בתור בעיית סיווג בינארית, שבה אנו קובעים ברמת הכתובת האם היא נכונה או לא, (2) ובתור בעיית דירוג (LTR), שבה אנו מדרגים עבור כל אדם את כלל הכתובות על מנת למצוא את הכתובת הנכונה. גישה זו מאפשרת לנו לבנות מודל שמבצע אופטימיזציה על מטריקות דירוג (למשל NDCG) ובכך להגיע לדיוק גבוה יותר

I B: AI in Medicine

Chair: Noa Dagan, *BGU and Clalit*

Patient Deterioration Prediction – Data Science and Statistical Analysis of a Prospective Tria

Yael Hershkovitz

Data Scientist and Biostatistician, I-NEXT-DATA, Tel Aviv Sourasky Medical Center

Early detection of clinical deterioration among hospitalized patients in internal medicine wards poses a significant challenge. This talk will present a machine learning model developed to predict deterioration events such as ICU admission, mechanical ventilation, or in-hospital mortality. In addition to an overview of the clinical context and the retrospective model development based on historical data, findings from a series of prospective trials conducted in several internal medicine departments will be shared.

These include a silent trial, a comparative study assessing the model's performance versus physicians' clinical judgment, and a live trial in which the model's output is presented to clinical teams via a dedicated application. The presentation will focus on the statistical analysis of these trials and the clinical implications of the model's predictions.

GenAI Applications for Doctors

Ben Gedassi

AI & GenAI Team Leader, Maccabi Healthcare Services

GenAI Applications for Doctors

A presentation of recent GenAI-based applications by Maccabi Healthcare Services, now integrated into various physician systems, designed to optimize workflow and improve the precision of medical care.

"Bat Kol", a GenAI-powered platform that transcribes doctor-patient encounters and generates clinical summaries and integrating them into the physician's operational systems—saving valuable time and reducing documentation burden.

Another application is a GenAI-based clinical assistant that provides physicians with real-time, chat-based access to thousands of clinical protocols directly within Maccabi's systems, enhancing decision-making at the point of care.

"See" the Model Through the OPTICA Lens: Responsible AI in Real-Life Practice

Stav Devons-Sberro, MD

Resident physician, Clinical AI & Research Department, Clalit Innovation

The integration of artificial intelligence (AI) into healthcare systems presents both opportunities and challenges, including variability in model performance across settings and the evolving nature of AI tools. Recognizing the need for structured evaluation, we developed the OPTICA (Organizational PerspecTIve Checklist for AI solutions adoption) – a framework designed to guide multidisciplinary teams through the responsible adoption of AI technologies within healthcare organizations. The checklist identifies key stakeholders and provides a clear sequence of questions: what should be asked, by whom, and how. It enables systematic assessment of ethical, clinical, and operational aspects across the AI lifecycle.

In this presentation, we will describe how the checklist has been implemented within Clalit and demonstrate its impact through real-world examples. These examples will focus on methodological gaps identified during AI evaluations, and how the framework helped address them. Our experience shows that the checklist not only supports decision-making, but also helps organizations recognize when AI tools are not ready for deployment, or what adaptations are required to implement them safely and effectively.

AI-ECG past, current, and future

Michal Cohen-Shelly

Head of AI ARC cardiovascular lab, Sheba Medical Center

Electrocardiograms (ECGs) are accessible, non-invasive diagnostic tools with growing potential in the era of artificial intelligence (AI). Leveraging deep learning (DL), we developed and validated multiple AI-ECG models to support early diagnosis and risk stratification of key cardiovascular conditions, including low ejection fraction (EF) and valvular heart disease, while also addressing the interpretability and calibration challenges crucial for clinical adoption.

To identify patients with reduced EF (≤40%), we evaluated several model architectures using a dataset of 50,814 ECGs from Sheba Medical Center. While various deep learning strategies were explored—including convolutional and transformer-based approaches—the convolutional models consistently demonstrated the highest performance, with an AUC of 0.89. A detailed calibration and error analysis revealed important demographic variations in model performance and emphasized the need for context-aware interpretation in clinical workflows.

To support broader cardiac screening, we also developed AI-ECG models for valvular pathologies including aortic stenosis (AS), mitral regurgitation (MR), and tricuspid regurgitation (TR). In parallel, we compared the diagnostic capacity of large language models (LLMs) such as GPT-4V for standard 12-lead ECG interpretation against traditional DL models, demonstrating both the promise and limitations of generative AI in ECG analysis. We explored multiple explainability techniques to better understand model decision-making and support clinical trust.

Future directions include a prospective validation study in an internal medicine department to assess real-world integration and performance, alongside retrospective analyses already conducted in the same clinical setting to evaluate the operational and diagnostic value of AI-ECG tools. This talk will highlight key research from Sheba Medical Center in the era of AI-ECG, while also reviewing global developments in research, commercialization, and clinical implementation.

II A: Applications of Statistics & AI in the Industry

Chair: Shimon Shahar, *TAU*

Forecasting water reserves in the Sea of Galilee

David Refaeli

The Statistics Lab of Tel-Aviv University

This talk outlines the development of a statistical modeling framework for forecasting the water level of Lake Kinneret (Sea of Galilee). The research, conducted by the Israel Water Authority in collaboration with Tel Aviv University's Center for AI and Data Science, was executed in two phases: (1) annual forecasting of available water resources using 32 years of historical rainfall and streamflow data; and (2) daily forecasting of water availability and Jordan River flow of 5 years using detailed meteorological and hydrological variables. Various modeling approaches were evaluated – including linear regression, time series models, additive models, etc. – with the final models outperforming existing benchmarks in most test years.

The results support improved water resource management under variable climatic conditions

Text analytics in medical notes with LLMs

Talia Tron

Briya

The healthcare sector accounts for approximately 30% of the world's data volume, positioning real-world data (RWD) as a "gold mine" within the medical field and a key driver for innovation in health-tech and the pharmaceutical industry. Yet, a substantial portion of essential patient information remains embedded in free-text formats such as physician notes, presenting significant challenges due to linguistic complexity, specialized terminology,

and inconsistent documentation practices. While recent advances in large language models (LLMs) have addressed some of these obstacles, the next generation of solutions demands a fundamental rethinking of data processing and analytics protocols. In this talk, I will present our approach to developing a foundational medical Hebrew LLM. We will discuss the unique integration between schema-driven question answering and knowledge augmentation, as well as strategies for incorporating medical expertise into both the training and evaluation phases.

Revisiting features selection for big data using randomness

Matityahu Sarafzadeh

PayPal

Feature selection is essential for building efficient machine learning models, yet for very large datasets it's often a complex and resource-intensive task. In this talk, I introduce the "Random Bar" method, a lightweight, intuitive technique rooted in early 20th-century statistical thinking that harnesses randomness to streamline modern ML workflows.

Unlike traditional approaches, which are often computationally expensive, detached from the model, or based on strong assumptions about the data, the method is model-agnostic, assumption-free, intuitive, and incredibly easy to apply. The Random Bar method for feature selection delivers the best of all worlds.

II B: Causal Inference & Design of Experiments

Chair: Micha Mandel, *HUJI*

Statistical Inference of Causality Between Opinions and Decisions with Ideas from Econometrics and Psychology

Yosi Rinott

The Hebrew University of Jerusalem

In the novel part of this talk, I will present our proposed statistical methodology for studying the causal relationship between people's opinions or knowledge about an issue (X) and their corresponding decisions or actions (Y). A simple example: X may represent an individual's opinion on the efficacy of a particular vaccine, while Y is the decision whether to get vaccinated.

Inferring causality in this context is challenging, especially because of the presence of unknown or unobserved confounders.

To address this, instrumental variables (IVs), a common tool in econometrics, are often employed. However, when examining the link between opinions and decisions, typical IV candidates may fail by acting as confounders themselves. We propose a novel IV designed for these scenarios.

The first part of the talk will provide the minimal required background on causal inference, covering the relevant concepts and terminology.

Joint work with Raanan Sulitzeanu-Kenan and Micha Mandel.

The Optimality of Pairwise Matching and Blocking Randomizations under Extreme Noise and Extreme Assignments

Adam Kapelner

Queens College, NY

We consider the general performance of the difference-in-means estimator in an equally-allocated two-arm randomized experiment under common experimental endpoints such as continuous (regression), incidence, proportion, count and uncensored survival. We consider two sources of randomness: the subject-specific assignments and the contribution of unobserved subject-specific measurements. We then examine mean squared error (MSE) performance under a new, more realistic simultaneous tail criterion. We prove that the pairwise matching design of Greevy et al. (2004) asymptotically performs best under the simultaneous tail criterion when compared with other blocking designs for all response types. We also prove that the optimal design must be less random than complete randomization and more random than any deterministic, optimized allocation. Theoretical results are supported by simulations.

Distinguishing Behavioral and Non-Behavioral Effects in Randomized Trials

Daniel Nevo

Tel Aviv University

In many randomized trials, treatment effects may arise not only from the intervention's intended mechanistic or biological action, but also from participants' beliefs and behavioral responses to perceived treatment assignment. This talk presents a framework for distinguishing these effects using standard trial data, under assumptions that are often plausible in practice. We show that if participant beliefs about treatment are measured, it is possible to define and estimate estimands that reflect either the non-behavioral effect of the treatment or the behavioral effect associated with perceived treatment status. While our motivating example involves vaccine trials, where behavioral changes following vaccination can influence infection risk, the approach applies more broadly. We demonstrate how conventional designs can be augmented with minimal changes to yield policy-relevant estimands, and illustrate the method using simulations based on data from an influenza vaccine trial.

The Israel four Cloud Seeding Experiment: Design and Analysis

David M Steinberg

Tel Aviv University

The Israel 4 Cloud Seeding Experiment was carried out over 7 years, from 2013 to 2020. This was a carefully designed randomized experiment to assess the effect of seeding on rainfall and provide a basis to evaluate its utility.
Operational seeding in northern Israel was initiated after two randomized experiments, the second ending in 1976, found a large and statistically significant effect of cloud seeding on rainfall. Observational studies in later years raised doubts as to the magnitude of the effect, possibly because of changing climatological conditions. The ensuing discussions convinced the Israel Water Authority (IWA) to run a current experiment to reassess the effect of seeding.

This talk will describe the design and analysis of the Israel 4 experiment.
Spoiler – the aftermath of Israel 4 was a decision by the IWA to terminate operational seeding.

The experiment was run under the auspices of the IWA and was supervised by an interdisciplinary steering committee. Yoav Benjamini and I provided statistical expertise as members of the steering committee.

A Convex Hull Volume-based Anomaly Detection Method

Uri Itai

GTIIT

Anomaly detection is a problem that appears in many industrial applications, both as a pre-process phase for data cleansing and as a stand-alone application. We present a novel anomaly detection method leveraging the convex hull (CH) structure of datasets. By observing how anomalies expand CH volume, our algorithm identifies out-of-distribution (OOD) data through iterative volume analysis. Evaluated across ten datasets, it performs on par with state-of-the-art methods while offering a simple, resource-efficient alternative. We also propose a criterion to predict when our method will outperform existing approaches.

מודים לחברות התומכות