



**VILNIUS
TECH**

Vilnius Gediminas
Technical University

The 13th International Conference

“Electrical, Electronic and Information Sciences – eStream 2026”

Abstract Proceedings

April 23–24, 2026
Vilnius, Lithuania

ISBN 978-609-476-464-6 [PDF]



**VILNIUS
TECH**

Vilnius Gediminas
Technical University

The 13th International Conference **“Electrical, Electronic and Information Sciences – eStream 2026”**

Abstract Proceedings

April 23–24, 2026
Vilnius, Lithuania

ISBN 978-609-476-464-6 (PDF)

Vilnius, 2026

13th International Conference “Electrical, Electronic and Information Sciences – eStream 2026”. Abstracts of presentations for the 13th International Conference “Electrical, Electronic and Information Sciences – eStream 2026”, Vilnius, Lithuania, April 23 to April 24, 2026. Vilnius Gediminas Technical University, 2026. 32 pages.

This book constitutes the abstracts of presentations presented for the 13th International Conference “Electrical, Electronic and Information Sciences – eStream 2026”, which is being held in Vilnius from April 23 to April 24, 2026 and organized by Vilnius Gediminas Technical University – VILNIUS TECH.

The conference is dedicated to fostering multidisciplinary research and the exchange of ideas across the rapidly evolving fields of Electrical and Electronic Engineering and Information Sciences.

The conference welcomes researchers, graduate students, and industry professionals from around the world to present their latest findings and innovations. The abstracts are printed without editing as presented by their authors.

For information contact:

Vilnius Gediminas Technical University – VILNIUS TECH,

Faculty of Electronics

Organizing Committee of International Conference eStream 2026

Plytinės g. 25, LT-10105 Vilnius, Lithuania

<https://conferences.vilniustech.lt/index.php/eStream>

VILNIUS TECH id 2026-052-K

ISBN 978-609-476-464-6 (PDF)

<https://doi.org/10.20334/2026-052-K>

Scientific Committee

Chair:

Prof. Dr. Dalius Navakas, Vice-rector for research and innovation,
VILNIUS TECH, Lithuania

Vice-Chair:

Prof. Dr. Vijitashwa Pandey, Dept. Chair and
Pawley Endowed Professor of Lean Studies, Oakland University, USA

Vice-Chair:

Prof. Dr. Darius Plonis, Head of Telecommunications Institute,
VILNIUS TECH, Lithuania

Host Representative:

Doc. Dr. Artūras Medeišis, Dean, Faculty of Electronics,
VILNIUS TECH, Lithuania

Organizing and Technical Program Committee

Chair:

Prof. Dr. Zigmantas Balevičius, Head of the Plasmonics and Nanophotonics Laboratory,
FTMC, Lithuania

Vice-Chair:

Prof. Dr. Anna Litvinenko, Elected professor and Deputy Director,
RTU AKITI, Latvia

Conference Secretary:

Dr. Justina Anulytė, Vice Dean for Research and Project Development,
Faculty of Electronics, VILNIUS TECH, Lithuania

Members/reviewers:

Aleksandr Vasjanov (VILNIUS TECH)
Algirdas Baškys (FTMC)
Andrius Katkevičius (VILNIUS TECH)
Andžej Borel (VILNIUS TECH)
Audrius Grainys (VILNIUS TECH)
Audrius Krukonis (VILNIUS TECH)
Dalia Mohammed Talat Ebrahim Ali (VILNIUS TECH)
Dainius Udris (VILNIUS TECH)
Dalius Matuzevičius (VILNIUS TECH)
Dalius Seliuta (FTMC)
Darius Guršnys (VILNIUS TECH)
Darius Miniotas (VILNIUS TECH)
Darius Plonis (VILNIUS TECH)
Diana Belova-Plonienė (VILNIUS TECH)
Dovilė Čibiraitė-Lukenskienė (FTMC)
Eldar Šabanovič (VILNIUS TECH)
Ervinas Gisleris (VILNIUS TECH)
Ernesta Bužavaitė-Vertelienė (FTMC)
Evaldas Stankevičius (VILNIUS TECH)
Gabriela Vdovjak (VILNIUS TECH)
Gediminas Gražulevičius (VILNIUS TECH)
Henrikas Giedra (VILNIUS TECH)
Inna Saiapina (National Technical University of Ukraine, KPI)
Yevgeniya Sulema (National Technical University of Ukraine, KPI)
Justina Anulytė (VILNIUS TECH)
Jaroslav Rokicki (VILNIUS TECH)
Leonid Kladovščikov (VILNIUS TECH)
Mantas Kazlauskas (VILNIUS TECH)
Marijan Jurgo (VILNIUS TECH)
Martynas Dumpis (VILNIUS TECH)
Martynas Šapurov (VILNIUS TECH)
Mindaugas Jankauskas (VILNIUS TECH)
Nerijus Paulauskas (VILNIUS TECH)
Pavel Stefanovič (VILNIUS TECH)
Raimondas Pomarnacki (VILNIUS TECH)
Robertas Ūselis (VILNIUS TECH)
Saulius Sakavičius (VILNIUS TECH)
Sonata Tolvaišienė (VILNIUS TECH)
Šarūnas Mikučionis (VILNIUS TECH)
Šarūnas Paulikas (VILNIUS TECH)
Tatjana Gric (VILNIUS TECH)
Tomyslav Sledevič (VILNIUS TECH)
Vaidotas Barzdėnas (VILNIUS TECH)
Vytautas Abromavičius (VILNIUS TECH)
Vytautas Urbanavičius (VILNIUS TECH)

MESSAGE FROM THE GENERAL CHAIR

Dear Colleagues, Professors, Scientists, and Students

We are very glad to welcome our colleagues – young scientists, researchers and practitioners to the 13th IEEE Open Conference of Electrical, Electronic and Information Sciences (eStream'2026), held in Vilnius Gediminas Technical University (VILNIUS TECH), Vilnius, Lithuania, on 23–24 April 2026. The eStream conferences aim to disseminate the research achievements between worldwide groups of scientists and engineers working in different areas of science to reach more tight relationships and generate new ideas for joint projects or other means of collaboration. We are proud to feature keynote presentations by distinguished scientists and industry experts, including:

Prof. Dr. Vilius Benetis, CEO of NRD Cyber Security, Lithuania

Assoc. Prof. Dr. Inna Saiapina, National Technical University of Ukraine, "Igor Sikorsky Kyiv Polytechnic Institute"

Prof. Dr. Vijitashwa Pandey, Dept. Chair and Pawley Endowed Professor of Lean Studies, Oakland University, USA

Dr. Gediminas Račiukaitis, Coordinator of ChipsC2 LT, State Research Institute Center for Physical Sciences and Technology, Lithuania

We also grateful to the organizers and speakers of the special session on standardization:

Amelie Leipprand, DIN (Germany)

Dr. Barbara Reiter, University of Graz (Austria)

Prof. Rudi Bekkers, Eindhoven University of Technology (Netherlands)

Dipl.-Ing. Robert Link, freesia Innovation e.U. (Austria)

Sokol Tominaj, Berlin Technical University (Germany)

The technical program of eStream'2026 was attended by 103 participants and contained 4 keynote speeches, 31 oral and 19 poster presentations delivered during 6 sessions. Each technical paper has been peer-reviewed by several independent referees of the international Technical Program Committee comprised of 44 members and, based on the recommendation of the reviewers, 17 papers are finally accepted for publication in the proceedings. We express our sincere thanks to the invited speaker and article authors, track chairs and reviewers who have made this conference a success.

Finally, we sincerely hope that you will find our conference and all contributions presented in the proceedings a valuable resource in your professional and research activities whether you are an academic, researcher, or a practicing professional!

Warm regards,

Conference Chair

Prof. Dr. Dalius Navakas,
Vice-rector for research and innovation,
VILNIUS TECH, Lithuania

CONTENTS

Parallelization of rainbow tables generation using message passing interface on homogenous nodes Mark VAINER	8
Time domain expansion and freezing investigation of fast analog signals using cascaded active analog memory cell array for applications in nuclear spectroscopy Vilius ŽALĖNAS	10
Geometric vs pixel representations for face shape classification Rūta STUNŽĖNAITĖ, Laura VENCKUTĖ, Artūras SERACKIS, Dalius MATUZEVIČIUS	11
Creating an autonomous driving agent using artificial intelligence methods Tomas BITVINSKAS, Dalius MATUZEVIČIUS	12
Defect detection strategies for glass bottle inspection Radoslav TITKO, Dalius MATUZEVIČIUS	13
Evaluating vision–language models for face shape classification Vytautas BUNEVIČIUS, Dalius MATUZEVIČIUS	15
An evaluation of image blending pipelines for generating realistic eyewear data Matsevytyi ANDRII, Artūras SERACKIS, Dalius MATUZEVIČIUS	17
Face shape classification for personalized eyewear recommendation Natavan JAFAROVA, Dalius MATUZEVIČIUS	18
Photogrammetric 3D reconstruction of asteroids: from synthetic baselines to Osiris-Rex mission data Vladas ŠATAS, Donatas NARBUTIS	19
Challenges of 3D reconstruction of small low-texture objects Arnoldas PUZANOVAS, Dalius MATUZEVIČIUS	20
Comparative analysis of high-voltage high-frequency pulse driver architectures Edgard ALEINIKOV, Vaidotas BARZDĖNAS	21
Image caption datasets: a review and a Lithuanian translation-based corpus Tautvydas KVIETKAUSKAS, Dmitrij KAMYŠANSKIS, Ūla ŠVEITERYTĖ	22
Wilkinson power divider 2.4 GHz frequency with distributed selective elements Kristina MILIAUSKAITĖ, Zigmas BALEVIČIUS	23
High-resolution patch-based semi-supervised learning for pancreatic cancer segmentation Istiaque AHMED, Shinthiya NOWSAIN PROMI, Kursat KOMURCU	24
Investigation of copper micro-path electrode design and line spacing on electrochromic device functionality Dainius BALKAUSKAS, Titas TAMOŠAUSKAS, Modestas SADAUSKAS, Edita VOITECHOVIČ, Romualdas TRUSOVAS, Karolis RATAUTAS, Alexandr BELOSLUDTSEV	26
Topology parameterization of impedance signals for spatial reconstruction and probe positioning in complex electrochemical domains Vadimas IVINSKIJ, Inga MORKVĖNAITĖ	27
Improving deepfake detection through better frame selection Mindaugas JANKAUSKAS, Artūras SERACKIS	28
Evaluation of learning-based algorithms for eyeglass reflection removal Henrikas GIEDRA, Dalius MATUZEVIČIUS	29
Investigation of drone light control and color calibration systems Lukas ALIŠAUSKAS	31

**The 13th International Conference
“Electrical, Electronic and Information Sciences – eStream 2026”**

Abstracts of poster presentations

PARALLELIZATION OF RAINBOW TABLES GENERATION USING MESSAGE PASSING INTERFACE ON HOMOGENOUS NODES

Mark VAINER*

Department of Graphical Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

*E-mail: mark.vainer@vilniustech.lt

Abstract. Password cracking in cryptanalysis and computer security involves recovering plaintext passwords from cryptographic hash values obtained through system compromise or data exfiltration. Because hash functions are one-way, attackers rely on indirect methods rather than direct inversion. Common approaches include brute-force attacks, which exhaustively test possible combinations but are computationally expensive, and dictionary attacks, which test passwords against predefined word lists and previously leaked credentials. While effective for weak passwords, these methods become inefficient or impractical when passwords are long and complex. Rainbow table attacks are more efficient. They rely on a large table of passwords and hash values to reveal the target passwords. However, while rainbow table attacks are considered fast and efficient, the pre-computation phase is a slow task to perform sequentially. Naturally, parallel computing is perceived as an obvious way to increase computational capabilities. However, the selection of the efficient parallel algorithm is highly dependable on the research field, the considered problem and the method used. In this work, the implementation of rainbow tables generation is proposed. The algorithm is written using the C++ programming language and the Message Passing Interface (MPI) library. The execution time of the proposed implementation was measured and speedup and efficiency were calculated. This approach modifies the classic master-slave parallelization model to reduce the inter-process communication overhead. The proposed method achieves a near-linear speedup and high parallel efficiency in the range of 95–99%.

```

BEGIN
  CALL MPI_Init()
  GET rank of current process
  GET size of total processes

  FOR i = 0 TO localM - 1 DO
    SP[i] = indexToString(...)

  FOR i = 0 TO localM - 1 DO
    SET hashed = hashPassword(SP[i])
    FOR j = 0 TO t - 1 DO
      SET reduced = reduceHash(hashed, j)
      SET hashed = hashPassword(reduced)
    EP[i] = reduced
  MASTER

  IF rank == 0 THEN
    FOR i = 0 TO localM - 1 DO
      INSERT local chain(i) into chains
    FOR p = 1 TO size - 1 DO
      MPI_RECEIVE chain from process p
      INSERT chain into chains
    SLAVES

  ELSE
    FOR i = 0 TO localM - 1 DO
      MPI_SEND chain(i) to process 0

  CALL MPI_Finalize()
END

```

Figure 1. Visual pseudocode for the load balancing implementation of the parallel rainbow table generation for homogenous resources

Table 1. Speedup for parallel large-size rainbow tables generation for SHA-256, SHA-512, MD5 and NTLMv2 cryptographic hash functions across an increasing number of computing nodes

Nodes count	SHA-256	SHA-512	MD5	NTLMv2
1	1	1	1	1
2	2.045	2.001	2.007	1.970
3	3.029	2.997	2.998	2.924
4	3.857	3.953	3.996	3.921
5	5.073	5.035	4.981	4.984
6	5.925	6.030	5.895	5.981
7	7.104	7.075	7.021	6.973
8	8.146	8.037	8.070	7.854
9	8.959	9.077	8.884	8.835
10	10.134	10.034	10.047	9.983
11	11.222	11.095	10.875	10.813
12	11.679	11.983	11.870	11.953
13	13.147	12.745	12.626	12.856
14	14.063	13.983	13.692	13.799
15	14.715	15.012	14.618	14.154




Figure 2. Execution time for parallel large-size rainbow tables generation for SHA-256, SHA-512, MD5 and NTLM (NTLMv2) cryptographic hash functions

Keywords: rainbow tables, password cracking, parallel computing, MPI master-slave model, cryptographic hash functions.

TIME DOMAIN EXPANSION AND FREEZING INVESTIGATION OF FAST ANALOG SIGNALS USING CASCADED ACTIVE ANALOG MEMORY CELL ARRAY FOR APPLICATIONS IN NUCLEAR SPECTROSCOPY

Vilius ŽALĖNAS*

Center for Physical Sciences and Technology, Savanorių pr. 231, LT-02300 Vilnius, Lithuania

**E-mail: vilius.zalenas@ftmc.lt*

Abstract. Nuclear spectroscopy has been closely related to fast analog signal processing and challenges that arise from it ever since its beginning. Luckily, today's scientific and technological achievements in semiconductors and integrated circuits allow us to use high-speed analog to digital converters and powerful processors that effectively eliminate the hardware and software signal and spectrum analysis problems. However, such systems are highly complicated from a signal integrity point of view, as well as the need for extensive and very specific knowledge and competence in particular embedded processor programming. Graphical user interface and data depiction via external software is another great challenge on its own, which overall results in very expensive, complicated and time-consuming to develop systems. This work is devoted to proposing a possible optimization to nuclear spectrometer hardware which would be equivalent to 100 MHz real-time sampling without using the actual 100 MHz analog to digital converter. While the converter is still an inevitable part of the proposed system, the signal would first be passed through an analog cell array with defined time offsets of the signal to sample moment. This approach could possibly allow us to capture 100 ns total duration scintillator detector signals (with rising edge times sometimes as low as 10 ns) in the form of several array cells DC voltages for later digitization. Since the sample time is only limited by the abundance of scintillation pulses and leakage currents in array cells, it could easily reach more than 100 μ s (possibly an improvement of three orders of magnitude). A possible compact module of this system could be used as a shield for the common Arduino, Raspberry Pi, ESP32 and other low-cost and user-friendly development boards to enable full-fledged nuclear spectroscopy and effectively replace complicated high-speed digital systems. This investigation developed an early-stage model of such a system and demonstrated its trustworthiness, need and value for further development.

Keywords: nuclear spectroscopy, particle energy analysis, signal processing.

GEOMETRIC VS PIXEL REPRESENTATIONS FOR FACE SHAPE CLASSIFICATION

Rūta STUNŽENAITĖ*, Laura VENCKUTĖ, Artūras SERACKIS, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

*E-mail: ruta.stunzenaite@stud.vilniustech.lt

Abstract. Automatic face shape classification remains challenging because of substantial intra-class variability and subtle geometric differences between classes. This study compares three feature representation strategies and four classical machine learning algorithms for multiclass recognition of five face shapes: heart, oblong, oval, round, and square. Experiments were conducted using a public dataset containing 1,000 facial images per class, with 800 images used for training and 200 for testing. The first representation used grayscale images resized to 64×64 pixels, flattened into 4,096-dimensional normalized raw-pixel vectors. The second employed MediaPipe Face Mesh to derive 13 manually engineered geometric descriptors, including facial proportions, jaw spread, and cheekbone-to-jawline angles. The third extracted all 468 facial landmarks and represented each face using 936 normalized x- and y-coordinate values. Landmark coordinates were centered relative to the nose tip and scaled to the $[-1, 1]$ interval to provide translation and scale invariance. Logistic Regression, Random Forest, Linear Support Vector Machine, and k-Nearest Neighbors classifiers were evaluated for each representation. The full landmark-based representation produced the strongest overall results, with the Linear Support Vector Machine achieving the highest test accuracy of 53%. These findings indicate that dense landmark coordinates preserve discriminative facial geometry more effectively than either raw pixel intensities or a limited set of manually designed geometric features. Nevertheless, the moderate accuracy and confusion patterns reveal considerable inter-class overlap. The oval class was particularly ambiguous and exhibited the lowest classification precision. Overall, the study demonstrates the value of structured landmark representations while highlighting the need for more robust feature-learning and class-separation strategies in face shape recognition.




Figure 1. Model accuracy




Figure 2. Confusion matrix for SVM model

Keywords: face shape classification, facial landmark detection, geometric feature extraction, MediaPipe Face Mesh, Support Vector Machine.

This project has received funding from the Research Council of Lithuania (LMTLT), agreement No S-ITP-24-12.

CREATING AN AUTONOMOUS DRIVING AGENT USING ARTIFICIAL INTELLIGENCE METHODS

Tomas BITVINSKAS^{1*}, Dalius MATUZEVIČIUS²

¹*Department of Information Technologies, Vilnius Gediminas Technical University, Vilnius, Lithuania*

²*Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania*

*E-mail: tomas.bitvinskas@stud.vilniustech.lt

Abstract. This work investigates an autonomous driving agent for closed-loop vehicle control in the BeamNG.tech simulator. The study evaluates whether a compact structured state representation and imitation learning can provide reliable lap completion on a selected test track without raw camera-based perception. Human driving demonstrations were converted into a 28-dimensional observation vector describing vehicle dynamics and local road geometry, including speed, yaw rate, acceleration, lateral offset, heading error, curvature, and look-ahead track features. A Behavioral Cloning (BC) policy with a Mixture Density Network (MDN) output head was trained to predict steering and longitudinal control, later mapped to throttle or brake commands. Four validated methods were compared in closed-loop evaluation at 30 Hz: Plain Behavioral Cloning (BC), Behavioral Cloning + Dataset Aggregation (BC + DAgger), Tuned Behavioral Cloning + Dataset Aggregation (BC + DAgger), and a model ensemble. The policies were evaluated using lap success rate, median lap time, and failure type. The ensemble achieved the best result, with a 97.7% mean success rate and a 139.4 s mean median lap time. Tuned BC + DAgger reached 62.0% success and 141.2 s, while BC + DAgger and Plain BC achieved 36.0% and 13.3% success, respectively. Remaining failures were mainly crashes and off-track events near curves and road boundaries. The results show that compact state-based behavioral cloning can produce strong simulated driving performance on the evaluated track, while unseen-track generalization and raw-vision perception remain future work.




Figure 1. Mean success rate and median lap time by method

Table 1. Compact summary of the main validated results

Method	Success rate	Median lap time
Ensemble	97.7%	139.4 s
Tuned BC + DAgger	62.0%	141.2 s
BC + DAgger	36.0%	142.9 s
Plain BC	13.3%	141.1 s

Keywords: autonomous driving, behavioral cloning, Mixture Density Network, BeamNG.tech, simulation.

DEFECT DETECTION STRATEGIES FOR GLASS BOTTLE INSPECTION

Radoslav TITKO*, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: radoslav.titko@stud.vilniustech.lt*

Abstract. Quality control remains a critical bottleneck in glass bottle manufacturing, where manual visual inspection typically detects only 70–80% of surface and structural defects. Although convolutional neural network (CNN)-based object detectors have demonstrated high accuracy in automated defect detection, most reported solutions rely on industrial-grade hardware that is often inaccessible to small and medium sized enterprises (SMEs). This study investigates whether modern deep learning architectures can achieve reliable defect detection under resource constrained, non-industrial conditions. Three object detection models: SSD MobileNetV2, Faster R-CNN (ResNet50), and YOLOv11m, were trained and evaluated on a dataset of 708 images, expanded to 1,848 through data augmentation, encompassing three defect classes: scratches, air bubbles, and chipped necks. Each model was trained both from scratch and using transfer learning with COCO pretrained weights on a consumer grade NVIDIA RTX 3060 (6 GB) GPU. YOLOv11m achieved the highest performance, reaching an mAP@0.5 of 0.9132 with transfer learning and 0.9071 when trained from scratch. It substantially outperformed SSD MobileNetV2 (0.7973) and Faster R-CNN (0.8140) under transfer learning while also delivering the fastest inference speed among the evaluated models (14.57 FPS). Transfer learning produced substantial accuracy improvements for the older architectures but only marginal gains for YOLOv11m, suggesting that modern detector architectures are less dependent on pretrained initialization. These findings demonstrate that state-of-the-art object detection can enable industrial grade defect inspection using limited training data and consumer grade hardware, providing a practical and cost-effective pathway toward automated quality control for SMEs and edge deployment scenarios.

Keywords: object detection, defect detection, YOLOv11, glass bottle inspection, CNN, quality control.




Figure 1. Types of defects: scratch, air bubble, and chipped neck

Table 1. Model performance comparison

Model	mAP@0.5	mAP@0.5-95
SSD MobileNetV2 (Scratch)	0.6654	0.4689
SSD MobileNetV2 (Transfer)	0.7973	0.5660
Faster R-CNN (Scratch)	0.5100	0.3384
Faster R-CNN (Transfer)	0.8140	0.5966
YOLOv11m (Scratch)	0.9071	0.7142
YOLOv11m (Transfer)	0.9132	0.7261




Figure 2. Detection results: 1 – YOLOv11, 2 – SSD, 3 – Faster R-CNN, 4 – Ground truth

References

- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778). <https://doi.org/10.1109/CVPR.2016.90>
- Khanam, R., & Hussain, M. (2024). YOLOv11: An overview of the key architectural enhancements. arXiv. <https://doi.org/10.48550/arXiv.2410.17725>
- Latina, M. A. E., Dela Cruz, J. V. R., & Delos Santos, F. D. (2022). Empty glass bottle defect detection based on deep learning with CNN using SSD MobileNetV2 model. In *2022 IEEE 14th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM)* (pp. 1–6). IEEE. <https://doi.org/10.1109/HNICEM57413.2022.10109368>
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft COCO: Common objects in context. In D. Fleet, T. Pajdla, B. Schiele, & T. Tuytelaars (Eds.), *Lecture notes in computer science: Vol. 8693. Computer Vision – ECCV 2014* (pp. 740–755). Springer. https://doi.org/10.1007/978-3-319-10602-1_48
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., & Berg, A. C. (2016). SSD: Single shot MultiBox detector. In B. Leibe, J. Matas, N. Sebe, & M. Welling (Eds.), *Lecture notes in computer science: Vol. 9905. Computer Vision – ECCV 2016* (pp. 21–37). Springer. https://doi.org/10.1007/978-3-319-46448-0_2
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28, 91–99. <https://doi.org/10.48550/arXiv.1506.01497>
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 4510–4520). IEEE. <https://doi.org/10.1109/CVPR.2018.00474>

EVALUATING VISION-LANGUAGE MODELS FOR FACE SHAPE CLASSIFICATION

Vytautas BUNEVIČIUS*, Dalius MATUZEVIČIUS

*Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania***E-mail: vytautas.bunevicius@stud.vilniustech.lt*

Abstract. Vision-language models (VLMs) achieve strong zero-shot performance across many multimodal tasks, yet their capacity for fine-grained geometric classification remains underexplored. Face shape classification is a representative fine-grained task and a core component of automated eyeglass frame recommendation, demanding sensitivity to subtle contour variation and global facial proportions rather than coarse object cues. This study systematically evaluates contemporary VLMs on five-class face shape classification (heart, oblong, oval, round, square) (niten19, 2019) under multiple inference and adaptation strategies.

Two 8B-parameter models, FaceLLM-8B and Qwen3-VL-8B, are assessed using zero-shot prompting, few-shot prompting, automated prompt optimization (DSPy MIPROv2 with Bayesian search and GPT-4.1 as teacher), and parameter-efficient fine-tuning. Across all prompt-based configurations, accuracy remains at 20–23%—random chance for five classes—with a pronounced single-class prediction bias. The failure of Bayesian prompt optimization to yield improvement indicates the limitation stems not from linguistic conditioning but from insufficient visual discrimination in the zero-shot regime.




Figure 1. Overview of the four evaluation paradigms




Figure 2. Accuracy comparison across all evaluated approaches

Fine-tuning Qwen3-VL-8B with QLoRA (4-bit quantization) on 3,500 labeled images—about 90 minutes on a single NVIDIA RTX 3090—raises accuracy to 85.9% (macro F1 = 0.856), demonstrating that VLMs can internalize discriminative geometric features under supervision. Recall remains uneven across classes, however; the oval class reaches only 65.5%, suggesting residual ambiguity between adjacent morphological categories or early overfitting.

These findings show that large-scale multimodal pretraining alone does not guarantee zero-shot transfer to fine-grained geometric tasks: explicit supervised adaptation is needed to extract task-relevant structural cues, with direct implications for deployable recommendation systems where reliability across visually similar categories is required. Ongoing work trains CNN baselines (ResNet, MobileNet) and evaluates contrastive VLMs (CLIP, SigLIP) to test whether contrastive pretraining confers improved geometric discrimination.




Figure 3. Confusion matrices: zero-shot collapse (left) vs. adapted model (right)

Keywords: vision–language models, face shape classification, fine-grained visual recognition, parameter-efficient fine-tuning, QLoRA.

This project has received funding from the Research Council of Lithuania (LMTLT), agreement No S-ITP-24-12.

References

nitens19. (2019). *Face shape dataset* [Data set]. Kaggle. <https://www.kaggle.com/datasets/nitens19/face-shape-dataset>

AN EVALUATION OF IMAGE BLENDING PIPELINES FOR GENERATING REALISTIC EYEWEAR DATA

Matsevytyi ANDRII*, Artūras SERACKIS, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: andrii.matsevytyi@vilniustech.lt*

Abstract. Realistic virtual try-on systems require visually coherent compositing of synthetic objects onto human subjects. In eyewear applications, the perceptual quality of the integration between 2D glasses and facial imagery directly determines user acceptance. A frequent artifact observed in naïve implementations is the so-called “pasted” appearance, characterized by inconsistent illumination, unrealistic transparency, and visible compositing boundaries. These artifacts are particularly problematic in real-time scenarios, where lighting conditions, camera exposure, and facial pose vary substantially. Simultaneously, the development of learning-based segmentation and generative pipelines for eyewear applications is hindered by the limited availability of annotated datasets. Manual data collection and pixel-level labeling for glasses segmentation are time-consuming and expensive. Simple overlay-based augmentation strategies degrade dataset realism and introduce biases that impair model generalization. The objective of this research was to systematically benchmark blending and harmonization algorithms in order to identify the most plausible methods for integrating 2D eyewear into facial images. Beyond perceptual realism, the long-term goal is the generation of a large-scale semi-synthetic dataset that alleviates the ground-truth bottleneck for training glasses segmentation models, including pipelines based on foundation models such as Segment Anything (SAM) combined with semi-supervised learning strategies. A total of 29 harmonization pipelines were investigated, implemented, and evaluated. The methods were grouped into three principal categories: (i) optimization-based approaches (7 methods), (ii) diffusion-based approaches (11 methods), and (iii) GAN-inspired approaches (11 methods). All methods were evaluated on a standardized benchmark set consisting of face images, eyewear templates, and precisely defined object masks. The optimization-based category included gradient-domain blending techniques such as Poisson blending, classical color transfer, histogram matching, and self-supervised pixel-level optimization procedures. These methods were formulated as constrained optimization problems in the image domain, aiming to minimize intensity discrepancies across compositing boundaries while preserving global color coherence. The diffusion-based category was constructed around Stable Diffusion-like generative models employing multi-modal conditioning. Conditioning signals included binary masks, structural edge guidance, and prompt-based semantic constraints. Additional variants incorporated latent space optimization, attention-weighted harmonization, and exemplar-driven style alignment. These pipelines were designed to refine composited images iteratively in the latent space, enabling structural consistency while adapting lighting, shading, and material properties. The GAN-inspired category consisted of mathematically implemented style-transfer and domain-translation frameworks inspired by architectures such as StyleGAN, CycleGAN, and StarGAN. Instead of training generative adversarial networks on dedicated datasets, structural principles of these architectures were analytically reproduced in NumPy-based formulations. This design choice enabled controlled analysis of their blending logic without confounding effects introduced by dataset bias. Evaluation was conducted exclusively on visual plausibility criteria. Three principal dimensions were assessed: (1) lighting and shading consistency, including realistic shadow casting on the nasal bridge and cheeks; (2) color consistency and transparency, ensuring preservation of frame color under ambient illumination and physically plausible lens refraction; and (3) boundary blending quality, measured by the absence of sharp compositing artifacts. The experimental analysis revealed clear performance stratification across categories. Optimization-based methods achieved satisfactory boundary blending and local gradient continuity. However, these approaches were unable to synthesize physically accurate drop shadows or simulate lens transparency, resulting in visually flat composites under complex lighting conditions. GAN-inspired methods demonstrated improved contrast alignment and skin-tone harmonization. In several cases, plausible global illumination adaptation was observed. Nevertheless, structural inconsistencies were occasionally introduced, including hallucinated frame components or unintended recoloring of eyewear to match facial chromatic distributions. Such behavior compromises the structural fidelity required for segmentation dataset generation. Diffusion-based methods with multi-modal conditioning substantially outperformed the remaining approaches across all evaluation dimensions. These pipelines generated coherent environmental reflections on lenses, realistic transparency effects, and plausible soft shadows consistent with facial geometry. Boundary transitions were smooth without suppressing structural details of the eyewear. The iterative latent refinement mechanism enabled adaptive harmonization under varying lighting conditions, indicating strong suitability for semi-synthetic data generation.

Keywords: virtual try-on harmonization, eyewear compositing, diffusion-based harmonization, optimization-based blending, GAN-inspired style transfer.

This project has received funding from the Research Council of Lithuania (LMTLT), agreement No S-ITP-24-12.

FACE SHAPE CLASSIFICATION FOR PERSONALIZED EYEWEAR RECOMMENDATION

Natavan JAFAROVA*, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

*E-mail: natavan.jafarova@stud.vilniustech.lt

Abstract. This research investigates face shape classification and its integration into an agentic artificial intelligence framework for personalized eyeglasses recommendation. A merged facial image dataset was constructed from two publicly available sources: the Kaggle Face Shape Dataset and the GitHub Face Shape Dataset. The final dataset contains 5,500 labeled images across five facial shape categories: Heart, Oblong, Oval, Round, and Square. This dataset was used to evaluate several convolutional and transformer-based architectures for multi-class face shape classification, including InceptionV3, InceptionV4, ResNet50, DenseNet121, Vision Transformer Small, and VGG16.

Among the evaluated models, InceptionV4 achieved the best performance, with a test accuracy of 75.27% and an F1-score of 80%. Confusion matrix analysis showed stable inter-class separation and relatively balanced misclassification patterns, suggesting satisfactory generalization across the five face shape categories. These results indicate that convolutional architectures remain effective for structured facial geometry classification, especially when class boundaries are visually subtle.

Beyond classification, the study implements an agentic AI recommendation pipeline. An uploaded facial image is processed to predict the face shape, which is then mapped to suitable eyeglass frame types using a structured compatibility schema. Candidate frames are retrieved from a curated catalog, ranked according to compatibility rules, and presented through explainable recommendations. The system also supports iterative refinement based on user preferences such as style and budget.

The study is further extended by comparing two recommendation paradigms: a perception-based CNN recommender and a knowledge-based LLM + RAG recommender. Both systems are evaluated using accuracy, precision, recall, F1-score, ranking quality, latency, inference time, GPU utilization, and API cost. Therefore, the proposed work contributes an empirical benchmark for face shape-driven eyewear recommendation and a systematic comparison of perception-centric and knowledge-centric agentic AI systems.

Keywords: face shape classification, personalized eyewear recommendation, convolutional neural networks, large language models, agentic AI.

References

- Deldjoo, Y., Nazary, F., Ramisa, A., McAuley, J., Pellegrini, G., Bellogin, A., & Di Noia, T. (2023). A review of modern fashion recommender systems. *ACM Computing Surveys*, 56(4), Article 87. <https://doi.org/10.1145/3624733>
- Jin, Y., Li, J., Gu, T., Liu, Y., Zhao, B., Lai, J., Gan, Z., Wang, Y., Wang, C., Tan, X., & Ma, L. (2025). Efficient multimodal large language models: A survey. *Visual Intelligence*, 3, Article 27. <https://doi.org/10.1007/s44267-025-00099-6>
- Khan, A., Sohail, A., Zahoor, U., & Qureshi, A. S. (2020). A survey of the recent architectures of deep convolutional neural networks. *Artificial Intelligence Review*, 53, 5455–5516. <https://doi.org/10.1007/s10462-020-09825-6>
- Rainio, O., Teuhio, J., & Klén, R. (2024). Evaluation metrics and statistical tests for machine learning. *Scientific Reports*, 14, Article 6086. <https://doi.org/10.1038/s41598-024-56706-x>
- Young, S., Natalia, F., Sudirman, S., & Ko, C. S. (2019). Eyeglasses frame selection based on oval face shape using convolutional neural network. *ICIC Express Letters, Part B: Applications*, 10(8), 707–715. <https://doi.org/10.24507/ICICELB.10.08.707>
- Zhang, Y., & Chen, X. (2020). Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval*, 14(1), 1–101. <https://doi.org/10.1561/15000000066>
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., Zhang, H., Gonzalez, J., & Stoica, I. (2023). Judging LLM-as-a-judge with MT-bench and chatbot arena. In *37th Conference on Neural Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks* (pp. 46595–46623). Neural Information Processing Systems Foundation. <https://doi.org/10.52202/075280-2020>

PHOTOGRAMMETRIC 3D RECONSTRUCTION OF ASTEROIDS: FROM SYNTHETIC BASELINES TO OSIRIS-REX MISSION DATA

Vladas ŠATAS*, Donatas NARBUTIS

Institute of Theoretical Physics and Astronomy, Faculty of Physics, Vilnius University, Vilnius, Lithuania

*E-mail: vladas.satas@ff.stud.vu.lt

Abstract. Optical sensor data serves as the primary information source during the early approach phases of space missions, preceding the operational range of active instruments like LIDAR. While photogrammetry offers a pathway for early shape and gravitational modelling, its application in space is constrained by radiometric and kinematic conditions. This research investigates the reliability and accuracy of Structure-from-Motion (SfM) and Multi-View Stereo (MVS) algorithms (Schönberger et al., 2016) under these specialized constraints. The study utilizes a three-phase progressive methodology. First, the computational pipeline was validated using procedurally generated synthetic asteroid model. Second, realistic space conditions were simulated, incorporating single-source illumination, high-contrast shadows, and restricted fly-by trajectories to stress-test the algorithms. Finally, an automated filtering and processing pipeline was applied to real optical imagery of asteroid (101955) Bennu, extracted from the OSIRIS-REx mission archive (Lauretta et al., 2017; NASA Planetary Data System, 2024). Reconstructed models were quantitatively evaluated against high-precision OLA LIDAR ground truth using the Cloud-to-Mesh (C2M) distance metric. Synthetic simulations demonstrated a trade-off in observational strategy, showing that maximizing image resolution yields higher geometric accuracy than simply increasing the density of captured frames during a fly-by. When applied to archival data, orbital phase sequence failed to reconstruct due to dynamic shadows and insufficient baseline overlap. However, approach phase sequences, where spacecraft’s linear motion combined with the asteroid’s rotation to create “virtual orbit” effect, yielded a successful global model. The final photogrammetric model achieved a mean geometric deviation of around 2 meters compared to the LIDAR reference, representing approximately 0.4% of the asteroid’s average diameter. The findings confirm that while monocular photogrammetry lacks the high-frequency topographic precision required for landing site selection, it successfully captures global volume and shape, making it a viable tool for early-phase mission planning and initial object characterization.

Keywords: photogrammetry, 3D reconstruction, structure-from-motion, multi-view stereo.

References

- Lauretta, D. S., Balram-Knutson, S. S., Beshore, E., Boynton, W. V., Drouet d’Aubigny, C., DellaGiustina, D. N., Enos, H. L., Gholish, D. R., Hergenrother, C. W., Howell, E. S., Johnson, C. A., Morton, E. T., Nolan, M. C., Rizk, B., Roper, H. L., Bartels, A. E., Bos, B. J., Dworkin, J. P., Highsmith, D. E., ... Sandford, S. A. (2017). OSIRIS-REx: Sample return from asteroid (101955) Bennu. *Space Science Reviews*, 212, 925–984. <https://arxiv.org/abs/1702.06981>
- NASA Planetary Data System. (2024). *Small bodies node: OSIRIS-REx data archive*. <https://sbn.psi.edu/pds/resource/orex/>
- Schönberger, J. L. et al. (2016). *COLMAP: Structure-from-Motion and Multi-View Stereo*. <https://colmap.github.io>

CHALLENGES OF 3D RECONSTRUCTION OF SMALL LOW-TEXTURE OBJECTS

Arnoldas PUZANOVAS*, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: arnoldas.puzanovas@vilniustech.lt*

Abstract. Small, low-texture objects remain challenging for structure-from-motion and multi-view stereo (SfM/MVS) reconstruction because they provide few stable local features, occupy limited image area, and may contain thin, glossy, or partially specular surfaces. Eyeglass frames represent a particularly demanding case due to their narrow geometry, small scale, low surface texture, and complex reflections. This study investigates 3D reconstruction as an input-conditioning problem, focusing on the acquisition, image-conversion, and scene-design choices required to preserve usable geometric information. A total of 34 image-acquisition datasets and 56 reconstruction attempts were produced using smartphones, compact cameras, and DSLR cameras. Images were acquired using either a turntable configuration or a moving-camera approach. Most datasets were captured under daylight, while selected datasets incorporated additional lighting or flash. Exposure parameters, including ISO, aperture, and shutter speed, were controlled where possible. Objects were placed on a white support surface or on one of four artificial patterned backgrounds.




Figure 1. Data set examples by surface variants




Figure 2. Reconstruction outcome by pipeline family

The image-generation and conversion pipeline was the dominant determinant of reconstruction feasibility. All DSLR native-JPG datasets produced 3D reconstructions, as did all compact-camera JPG datasets, although compact-camera reconstruction quality was more sensitive to image count, ISO, and support-surface design. Standard smartphone computational photography frequently degraded or removed geometry-relevant image cues. Structured support surfaces increased feature density and SfM stability, enabling successful reconstruction from fewer images; equalized random-tile patterns performed best. The results demonstrate that pipeline compatibility is more important than image count alone. Once a reconstructable image pipeline is established, structured scene design substantially improves feature extraction and reconstruction reliability. Within the evaluated workflow, JPG was the most stable format. Acquisition protocols for small, low-texture objects should therefore jointly optimize image formation, conversion, exposure, and support-surface structure.

Keywords: 3D reconstruction, eyeglasses reconstruction, structure-from-motion, multi-view stereo, low-texture objects.

This project was supported by the Research Council of Lithuania (LMTLT), agreement No S-ITP-24-12.

COMPARATIVE ANALYSIS OF HIGH-VOLTAGE HIGH-FREQUENCY PULSE DRIVER ARCHITECTURES

Edgard ALEINIKOV*, Vaidotas BARZDĖNAS

*Department of Computer Science and Communications Technologies,
Vilnius Gediminas Technical University, Vilnius, Lithuania*

**E-mail: edgard.aleinikov@vilniustech.lt*

Abstract. This article presents a comparative analysis of high-voltage, high-frequency pulse generation techniques for applications requiring nanosecond pulses up to kilovolt amplitudes. Key design challenges are discussed with particular attention to thermal management using air, liquid, solid-state, and phase-change cooling methods. Common pulse generation architectures including vacuum tubes, voltage multipliers, Marx generators, Blumlein structures, pulse-forming networks, Tesla transformers, switching-mode power supplies, solid-state switches, and high-voltage operational amplifiers are examined in terms of cost, complexity, stability and suitability for driving capacitive loads. Hybrid solutions that combine solid-state switching with modular multipliers or pulse-forming circuits are shown to offer a favorable balance between efficiency, compactness, and reliability. These results provide practical guidance for the development of next-generation high-performance high-voltage pulse drivers.

Keywords: architectural comparison, fast pulse rise time, high-frequency, high-voltage, Marx generator, performance analysis, pulse driver technologies, pulse generation techniques, pulse generator, solid-state switches, switching-mode power supply, thermal management.

IMAGE CAPTION DATASETS: A REVIEW AND A LITHUANIAN TRANSLATION-BASED CORPUS

Tautvydas KVIETKAUSKAS*, Dmitrij KAMYŠANSKIS, Ūla ŠVEITERYTĖ

Faculty of Fundamental Science, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: tautvydas.kvietkauskas@vilniustech.lt*

Abstract. Data are a key factor in the development of accurate and reliable artificial intelligence models. Due to the popularity of large language models, a lot of attention has been focused on text data. Over time, the use of image and text has increased, which has led to the development of more complex vision-language models. Image-caption tasks require not only suitable images, but also appropriate descriptions. Although there exist many opensource datasets, their reliability and validity are not always clear. Consequently, data may contain errors, whether described by humans or artificial intelligence. However, compared to ChatGPT or other text generators, humans are considered experts. Datasets of image-caption type can be used to create or fine-tune vision-language models. During the study, it was observed that the datasets were not available in Lithuanian. Therefore, a new corpus of images and Lithuanian captions was created using open datasets annotated by experts rather than generators. The results of this research may be useful to other researchers in understanding the scope of vision-language datasets. The translated Lithuanian dataset is available on request.

Keywords: image-caption, vision-language models, dataset, Lithuanian, translation, data preprocessing.

WILKINSON POWER DIVIDER 2.4 GHZ FREQUENCY WITH DISTRIBUTED SELECTIVE ELEMENTS

Kristina MILIAUSKAITE*, Zigmas BALEVIČIUS

*Department of Electronic Systems, Faculty of Electronics, Vilnius Gediminas Technical University,
Vilnius, Lithuania*

*E-mail: k.miliauskaite@stud.vilniustech.lt

Abstract. Power dividers are key components in contemporary radio frequency (RF) and microwave systems, primarily used to distribute input signal power evenly among multiple output ports. Among the different power divider configurations, the Wilkinson power divider is especially popular due to its advantageous properties, such as equal power splitting and high isolation between output ports (Zafar et al., 2022). Although this topology inherently provides strong isolation, a distributed-element bandpass filter is placed at the input of the Wilkinson divider to further enhance signal quality and enable more precise bandwidth control.

The power divider with bandpass filters is designed using microstrips. While designing the circuit low-loss dielectric materials are used, microstrips can provide excellent performance with minimal signal attenuation and good impedance control at RF and microwave frequencies. Low-loss dielectric material Rogers 4350B is used, which has dielectric constant $\epsilon_r = 3.66$. Keysight Advanced Design System (ADS) software is used for circuit design and simulation (Keysight, 2024). When simulating the reflection coefficient S_{11} , the center frequency of 2.4 GHz reaches a reflection of approximately -30 dB, which means that only a small portion of the signal is reflected back to the source. This indicates good impedance matching at the operating frequency and efficient power transfer to the load. The forward transmission coefficient S_{21} is close to -3 dB, demonstrating nearly equal power division between the two output ports.

Integrating a distributed-element bandpass filter with a Wilkinson power divider provides an efficient solution for achieving both signal filtering and power division within the 2.4 GHz band, while preserving good impedance matching and high isolation between output ports. Such circuits can serve as a basis for further development of filtering power dividers for use in compact RF circuit which operate in the 2.4 GHz band, such as Wi-Fi and Bluetooth transceivers, multi-antenna access points or for measurement/testing setups where a single RF signal must be split into two or more.

Keywords: Wilkinson power divider, bandpass filter, microstrip, 2.4 GHz, S-parameters.

References

- Keysight. (2024). *Advanced Design System (ADS) circuit design cookbook 2.0*. <https://www.keysight.com/us/en/assets/7018-04800/application-notes/5992-0707.pdf>
- Zafar, M., Ekpo, S., George, J., Sheedy, P., Uko, M., & Gibson, A. (2022). Hybrid power divider and combiner for passive RFID tag wireless energy harvesting. *IEEE Access*, 10, 502–515. <https://doi.org/10.1109/ACCESS.2021.3138070>
- Zhu, J., Wang, Y., Zhao, K., Wang, Y., Fu, C., & Man, K. (2024). Assessment of measurement uncertainty for S-parameter measurement based on covariance matrix. *Sensors*, 24(11), Article 3668. <https://doi.org/10.3390/s24113668>

HIGH-RESOLUTION PATCH-BASED SEMI-SUPERVISED LEARNING FOR PANCREATIC CANCER SEGMENTATION

Istiaque AHMED, Shinthiya NOWSAIN PROMI*, Kursat KOMURCU

Information Systems Software Engineering Studies, Vilnius Gediminas Technical University,
Vilnius, Lithuania

*E-mail: shinthiya-nowsain.promi@stud.vilniustech.lt

Abstract. Precise segmentation of the pancreas from CT scans is important for diagnosis but remains difficult due to high anatomical variability and limited annotations. Conventional methods typically down sample inputs to fit GPU memory, which destroys fine-grained boundary information. This study proposes a high-resolution patch-based framework that preserves native CT resolution (512×512) through 256×256 patch extraction. A lightweight 2D U-Net (7.8M parameters) achieves 0.809 ± 0.015 Dice across four independent runs (peak 0.849) on MSD Task07 Pancreas, competitive with nnU-Net (0.823, 88M parameters) and clearly above SAM (0.705). Fourier analysis shows that downsampling removes 38% of boundary edge information. To address annotation scarcity, the framework is combined with Uncertainty-Aware Mean Teacher (UA-MT), which recovers 99.3% of supervised performance using only 50% of labelled data. Zero-shot evaluation on 80 external TCIA volumes shows that UA-MT (0.603 Dice) generalises better than the fully supervised model (0.544) across domains. The backbone is a 2D architecture operating on individual axial slices, so we report cross-slice consistency rather than volumetric modelling: across anatomical depths within a volume, the patch-based model stays stable, whereas heavier 3D and transformer baselines collapse on boundary slices. It also runs on a single GPU with far fewer parameters (7.8M vs 88M for nnU-Net) and needs no architecture search. These results indicate that, for the pancreas, preserving native voxel resolution matters more than capturing whole-abdomen context, and that uncertainty-weighted consistency regularisation is well suited to its ambiguous boundaries (Figure 1, Table 1).




Figure 1. Qualitative comparison of pancreas segmentation

Table 1. Quantitative comparison on MSD Task07 Pancreas (Dice, higher is better)

Method	Input res.	Labelled	Params (M)	Dice
nnU-Net	Auto	100%	88.0	0.823
SAM (box)	Fixed	100%	91.0	0.705
Ablation (downsampled)	256×256	100%	7.8	0.311
Patch U-Net (ours, best)	256×256 patches	100%	7.8	0.849
UA-MT (ours)	256×256 patches	50%	7.8	0.803

Keywords: pancreas segmentation, semi-supervised learning, patch-based deep learning, Mean Teacher, CT imaging.

References

- Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B. A., Litjens, G., Menze, B., Ronneberger, O., Summers, R. M., van Ginneken, B., Bilello, M., Bilic, P., Christ, P. F., Do, R. K. G., Gollub, M. J., Heckers, S. H., Huisman, H., Jarnagin, W. R., ... Cardoso, M. J. (2022). The medical segmentation decathlon. *Nature Communications*, 13, Article 4128. <https://doi.org/10.1038/s41467-022-30695-9>
- Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18, 203–211. <https://doi.org/10.1038/s41592-020-01008-z>
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., & Girshick, R. (2023). Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 4015–4026). IEEE. <https://doi.org/10.1109/ICCV51070.2023.00371>

- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells & A. F. Frangi (Eds.), *Medical image computing and computer-assisted intervention (MICCAI): Vol. 9351* (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28
- Yu, L., Wang, S., Li, X., Fu, C.-W., & Heng, P.-A. (2019). Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation. In D. Shen, T. Liu, L. H. Staib, C. Essert, S. Zhou, P.-T. Yap, & A. Khan (Eds.), *Medical image computing and computer-assisted intervention (MICCAI): Vol. 11765* (pp. 605–613). Springer. https://doi.org/10.1007/978-3-030-32245-8_67

INVESTIGATION OF COPPER MICRO-PATH ELECTRODE DESIGN AND LINE SPACING ON ELECTROCHROMIC DEVICE FUNCTIONALITY

Dainius BALKAUSKAS*, Titas TAMOŠAUSKAS, Modestas SADAUSKAS, Edita VOITECHOVIČ, Romualdas TRUSOVAS, Karolis RATAUTAS, Alexandr BELOSLUDTSEV

*Center for Physical Sciences and Technology, Department of Laser Technologies,
Savanorių pr. 231, Vilnius 02300, Lithuania*

**E-mail: dainius.balkauskas@ftmc.lt*

Abstract. Electrochromic devices (ECDs) are increasingly important for applications such as smart windows and adaptive optical systems, yet their performance and cost remain strongly dependent on the choice of transparent electrodes. This study addresses the challenge of replacing conventional indium-based conductors with low-cost copper while maintaining optical transparency, electrochemical stability, and device durability. The objective was to evaluate how copper line-pattern electrodes of varying geometries perform in ECDs employing PEDOT:PSS or WO₃ electrochromic layers and LiClO₄ or H₂SO₄ electrolytes. Transparent microscale copper electrodes were fabricated using the SSAIL technique (Balkauskas et al., 2025), and devices were assembled and characterized through optical transmittance measurements (Sadauskas et al., 2025), switching-speed analysis, and cycling stability tests. The results show that electrode geometry strongly influences optical modulation and switching kinetics, with narrower line spacing yielding higher transmittance but increased resistance. Electrolyte choice proved critical: PEDOT:PSS devices with LiClO₄ demonstrated stable operation with copper electrodes, while H₂SO₄ caused rapid copper degradation and loss of electrochromic functionality. These findings highlight the importance of electrode–electrolyte compatibility and provide practical guidelines for designing durable, cost-effective ECDs using copper-based transparent electrodes.

Keywords: electrochromic device.

References

- Balkauskas, D., Kvietkauskas, E., Sadauskas, M., Vrubliauskaitė, V., Trusovas, R., Ratautas, K., & Belosludtsev, A. (2025). Transparent microscale electrodes formed by laser-induced metal deposition on glass for smart windows application. *Optics & Laser Technology*, 192, Article 113669. <https://doi.org/10.1016/j.optlastec.2025.113669>
- Sadauskas, M., Vrubliauskaitė, V., Kvietkauskas, E., Stankevičienė, I., Balkauskas, D., Trusovas, R., Belosludtsev, A., & Ratautas, K. (2025). Laser-induced microscale copper trace deposition on glass and PET substrates. *Optics & Laser Technology*, 183, Article 112355. <https://doi.org/10.1016/j.optlastec.2024.112355>

TOPOLOGY PARAMETERIZATION OF IMPEDANCE SIGNALS FOR SPATIAL RECONSTRUCTION AND PROBE POSITIONING IN COMPLEX ELECTROCHEMICAL DOMAINS

Vadimas IVINSKIJ*, Inga MORKVĖNAITĖ

Department of Electrical Engineering, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: vadimas.ivinskij@vilniustech.lt*

Abstract. Electrochemical impedance spectroscopy (EIS) is a widely used method for characterizing electrochemical systems because it captures information about charge transfer, interfacial behavior, and diffusion over a range of frequencies. In spatially resolved techniques such as scanning electrochemical microscopy (SECM), however, impedance measurements are usually sparse, locally defined, and incomplete. As a result, the measured spectra do not directly provide a global map of the electrochemical structure of the scanned domain. This limitation becomes especially important when the objective is not only to interpret the sample but also to guide probe motion toward the most informative regions during scanning. This study proposes a topology-aware framework for reconstructing electrochemical structure from sparse impedance spectra and using the reconstructed structure to guide autonomous probe navigation. The central idea is to treat each impedance spectrum, written as $Z(x, \omega)$, as a spatial signal linked to a location x in the scanning domain Ω . Instead of analyzing each spectrum independently, the method transforms the measured spectra into compact feature representations, denoted by $\phi(Z)$, that summarize their spectral and electrochemical characteristics. These features are then used to reconstruct a continuous electrochemical activity field $\hat{u}(x)$, which serves as a spatial representation of electrochemical relevance across the domain. To perform this reconstruction, the study adopts a Neural Tangent Kernel (NTK) approximation. In this framework, the learned field $\hat{u}(x)$ is reconstructed from sparse measured locations X_{train} by using a kernel-based formulation that captures similarity between impedance-derived features. The advantage of this approach is that it provides a smooth and stable interpolation of the electrochemical activity field even when only a limited fraction of spatial points has been measured. In addition to the reconstructed field, the methodology also estimates an uncertainty field $\sigma(x)$, which identifies where the prediction is strongly supported by available measurements and where it depends more heavily on extrapolation. This distinction is important because it allows the framework to separate locally reliable reconstruction from less certain globally inferred regions. Once the electrochemical field has been reconstructed, biased topology analysis is applied in order to identify meaningful spatial structures. This is done by defining level sets X_τ , which correspond to all points in the domain for which the reconstructed field $\hat{u}(x)$ exceeds a chosen threshold τ . By varying τ , the method tracks structural changes in the field, such as the appearance, disappearance, or merging of connected regions, as well as the formation of holes. These features are interpreted as electrochemically meaningful structures, including active sample regions, defect cavities, and diffusion information saturated transition zones. In this way, the field reconstruction is not evaluated only in terms of point wise numerical accuracy, but also in terms of its ability to preserve structural organization. The topological information is then incorporated into a navigation cost function $c(x)$, which combines electrochemical relevance, uncertainty, and safety. Locations with high electrochemical activity $\hat{u}(x)$ are treated as attractive scanning targets, while locations with high uncertainty $\sigma(x)$ or low safety distance $d_{\text{safe}}(x)$ are treated more cautiously. This cost field is used within a Rapidly Exploring Random Tree (RRT) planner to generate probe trajectories $p(t)$. Unlike a purely geometric planner, the proposed navigation strategy does not explore the domain blindly. Instead, it guides the probe toward electrochemically relevant and structurally meaningful regions inferred directly from the measured impedance spectra. The framework was evaluated on a synthetic electrochemical domain containing active regions, defects, and diffusion-dominated areas. Spatially varying equivalent-circuit parameters were used to generate impedance spectra over a logarithmic frequency range, and sparse sampling was then applied to simulate realistic scanning conditions. The results showed that the NTK-based reconstruction recovered the dominant structure of the electrochemical field with good spatial coherence, while the associated uncertainty map correctly identified regions of lower confidence. The threshold topology analysis preserved the principal connected components and hole structures of the original domain. In the navigation stage, the topology navigation planner generated more targeted probe paths than baseline RRT strategies, directing the probe toward informative electrochemical regions while avoiding less relevant or uncertain areas. Overall, the study indicates that sparse impedance measurements can be transformed into a topology navigation aware spatial representation that supports both electrochemical interpretation and autonomous probe guidance. The proposed framework therefore contributes to the development of intelligent electrochemical sensing systems in which impedance-based reconstruction and navigation are directly coupled within a single data driven pipeline.

Keywords: machine learning, feature engineering, MLP, FFT mapping, gradient flow dynamics, FF, NTK neural networks.

IMPROVING DEEFAKE DETECTION THROUGH BETTER FRAME SELECTION

Mindaugas JANKAUSKAS*, Artūras SERACKIS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

*E-mail: m.jankauskas@vilniustech.lt

Abstract. Video deepfake detectors often classify a video from only a limited subset of frames, yet the number and selection of those frames directly influence both detection accuracy and inference cost. This study evaluates inference-time frame selection as an explicit experimental factor in video deepfake detection. We compare 12 selection strategies, including uniform, quality-based, motion-based, shot-aware, face-utility, and landmark-aware methods, across frame budgets of 2, 4, 8, 16, and 32 frames per video. Experiments use 300 videos from Celeb-DF-v2, Celeb-DF++, and FaceForensics++ (Rössler et al., 2019; Li et al., 2020), with four validated detectors: SBI, FSBI, GenConViT, and GenD. To keep the comparison controlled, the selected frames are materialized in a shared selection cache and reused across detectors. Results show that 32 frames achieve the highest overall mean AUC (0.9257), but smaller budgets remain competitive: compared with 32 frames, the AUC drop is only 0.0039 at 16 frames and 0.0069 at 8 frames. GenD obtains the strongest mean AUC across the tested budgets, while FSBI provides the fastest inference at the largest budget. Runtime analysis on an NVIDIA H100 server further shows that more complex selection strategies can substantially increase selection cost without consistently improving AUC. These findings suggest that practical video deepfake detection systems should balance accuracy and computational efficiency by treating frame selection and frame budget as first-class design choices.




Figure 1. Mean AUC across frame budgets with detector-wise loss relative to 32 frames

Keywords: video deepfake detection, frame selection, frame budget, inference-time sampling, runtime trade-off.

References

- Li, Y., Yang, X., Sun, P., Qi, H., & Lyu, S. (2020). Celeb-DF: A large-scale challenging dataset for deepfake forensics. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 3207–3216). IEEE.
<https://doi.org/10.1109/CVPR42600.2020.00327>
- Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., & Niessner, M. (2019). FaceForensics++: Learning to detect manipulated facial images. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)* (pp. 1–11). IEEE.
<https://doi.org/10.1109/ICCV.2019.00009>

EVALUATION OF LEARNING-BASED ALGORITHMS FOR EYEGLASS REFLECTION REMOVAL

Henrikas GIEDRA*, Dalius MATUZEVIČIUS

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: henrikas.giedra@vilniustech.lt*

Abstract. Reflections on eyeglass lenses can obscure the eye region in facial images and reduce the performance of eye-geometry-related tasks. While recent reflection-removal methods have improved visual restoration quality (Wan et al., 2021, 2023; Zou et al., 2024), less attention has been given to their effect on eye-region geometry. This study evaluates whether learning-based eyeglass reflection removal improves both visual appearance and downstream iris-center localization accuracy.

The evaluation used paired synthetic and real image data (Figure 1). Synthetic samples were generated using diffusion models and real eyeglass-frame assets, while real samples were acquired under fixed conditions using a mannequin head with eyeglasses. Five reflection-removal methods from different model families were compared: WindowSeat, Dereflexion Any Image, DSIT, DSRNet, and RDNet. A modular lens-cutout pipeline was also applied: eyeglass lenses were segmented using Meta SAM2, processed independently, and composited back into the original image. Iris centers were detected using MediaPipe Face Landmarker. Localization accuracy was measured using Normalized Iris-center Distance (NID), which calculates the Euclidean distance between the detected iris center and the clean reference iris center, normalized by the reference iris diameter.




Figure 1. Examples of synthetic and real evaluation data

The results show that reflection removal can improve iris-center localization in both synthetic and real images. Recovery was computed as the relative reduction in NID after reflection removal. Dereflexion Any Image achieved the highest average recovery on synthetic data, reaching 38.1%, while WindowSeat performed best on real mannequin images, reaching 37.0% recovery (Table 1).

Table 1. NID recovery (%). Synthetic values are averaged across reflection types

Model	Synthetic (%)	Real (%)
Dereflexion Any Image	38.1	23.7
DSIT	29.9	17.2
DSRNet	22.8	19.5
RDNet	28.7	27.7
WindowSeat	32.7	37.0

Qualitative examples also show improved iris visibility and more stable iris-center estimates after reflection removal (Figure 2).

Overall, the results indicate that eyeglass reflection removal should be evaluated using task-oriented metrics, since visually improved outputs do not always guarantee geometric accuracy. Future work will focus on dedicated eyeglass-lens segmentation and fine-tuning reflection-removal models on synthetic paired training data.




Figure 2. Qualitative examples from synthetic (top) and real mannequin (bottom) data

Keywords: computer vision, deep learning, reflection removal, eye-region geometry, synthetic image data.

This project was supported by the Research Council of Lithuania (LMTLT), agreement No S-ITP-24-12.

References

- Wan, R., Shi, B., Li, H., Duan, L.-Y., & Kot, A. C. (2021). Face image reflection removal. *International Journal of Computer Vision*, 129(2), 385–399. <https://doi.org/10.1007/s11263-020-01372-5>
- Wan, R., Shi, B., Li, H., Hong, Y., Duan, L.-Y., & Kot, A. C. (2023). Benchmarking single-image reflection removal algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2), 1424–1441. <https://doi.org/10.1109/TPAMI.2022.3168560>
- Zou, W., Lu, X., Yi, Z., Zhang, L., Fu, G., Li, P., & Xiao, C. (2024). Eyeglass reflection removal with joint learning of reflection elimination and content inpainting. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(10), 10266–10280. <https://doi.org/10.1109/TCSVT.2024.3405576>

INVESTIGATION OF DRONE LIGHT CONTROL AND COLOR CALIBRATION SYSTEMS

Lukas ALIŠAUSKAS*

Department of Electronic Systems, Vilnius Gediminas Technical University, Vilnius, Lithuania

**E-mail: lukas.alisauskas@stud.vilniustech.lt*

Abstract. Drone lighting systems with RGB and RGBW LEDs are used in applications where it is important that all drones produce similar brightness and color. This is especially important in multi-drone shows and signaling systems, where even small differences between drones can be clearly visible. In practice, such differences may appear due to LED manufacturing variations, temperature changes, and aging. Therefore, it is important to investigate suitable light control and color calibration methods for drone lighting systems.

This work studies practical methods for LED control and color calibration in drone applications. It considers common control methods, such as pulse-width modulation and current-based control, as well as calibration methods such as visual adjustment, matrix-based correction, and lookup-table correction. The aim is to identify a practical and reliable calibration method that can be applied in real drone systems.

The proposed solution is based on an external calibration stand. In this approach, a drone is placed in a controlled measurement setup, test light signals are generated, and the emitted light is measured using an optical sensor. Based on these measurements, correction coefficients are calculated and stored in the drone controller for use until the next recalibration cycle. The possibility of using temperature measurements for additional correction is also considered.

Initial experiments have already been carried out. The results showed that separately assembled LED lighting units can produce visible differences even when controlled in the same way. The measurements also showed that the selected setup can provide stable and repeatable readings, and that preliminary correction improves the similarity of the emitted light between the tested units.

The described approach offers a practical way to improve color consistency and repeatability in drone lighting systems during both production and maintenance. It also provides a useful basis for further development of drone light calibration methods.

Keywords: drone lighting calibration, LED brightness and color consistency, pulse-width and current-based control, optical sensor-based calibration stand, lookup-table, coefficient correction.

13th International Conference
“Electrical, Electronic and Information Sciences – eStream 2026”.
Abstract Proceedings

Apimtis 8,0 sp. l.
Leidinio el. versija <https://doi.org/10.20334/2026-052-K>
Vilniaus Gedimino technikos universitetas
Saulėtekio al. 11, 10223 Vilnius