45th Symposium on Information Theory and Signal Processing (SITB 2025)

Preliminary proceedings

Two-stage Bayesian Non-parametric Clustering for Channel with Intersymbol Interference and Bursty Impulsive Noise

Chin-Hung Chen[⋆], Boris Karanov[†], Yan Wu[‡], and Wim van Houtum[⋆‡]

⋆ Information and Communication Theory Lab, Eindhoven University of Technology, The Netherlands
† Communications Engineering Lab, Karlsruhe Institute of Technology, Germany
‡ NXP Semiconductors, Eindhoven, The Netherlands

In wireless communication environments, intersymbol interference (ISI) is a well-known issue often caused by multipath propagation. In addition, in recent years, a new type of interference has emerged, characterized by consecutive high-power spikes known as bursty impulsive noise (IN). This interference has become prevalent due to the increased use of power devices, which severely impaired the reception quality of high-sensitivity receivers [1]. In optimal trellis-based receiver designs that aim for robust channel estimation [2, 3] and symbol detection [4], knowing the number of channel states is often considered essential. However, in practice, this information is frequently unavailable. In this study, we propose a two-stage Bayesian nonparametric method for precisely estimating the number of states of channels affected by ISI and bursty IN. The first stage uses a Bayesian non-parametric collapse Gibbs sampling technique to determine the number of ISI states. After identifying the dominant clusters in the first stage, this information is passed on to the second stage. This stage utilizes a Bayesian hidden Markov model architecture to determine the optimal number of IN states based on the scores of the model evidence.

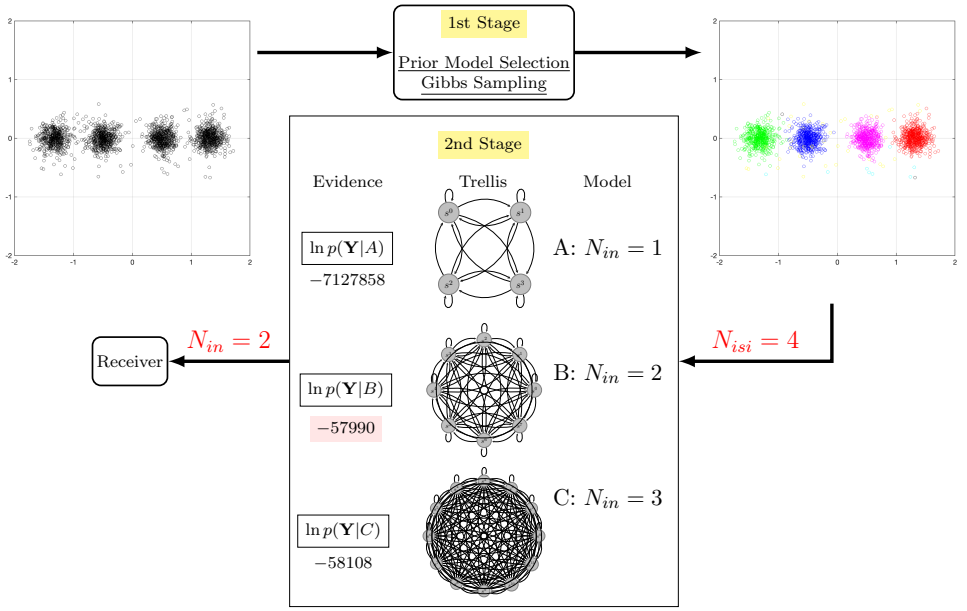

Figure 1: Diagram of the two-stage state provisioning algorithm.

[1] C.-H. Chen, W.-H. Huang, B. Karanov, Y. Wu, A. Young, W. van Houtum, "Analysis of impulsive interference in digital audio broadcasting systems in electric vehicles," *Symposium on Information Theory and Signal Processing in the Benelux (SITB)*, 2024.
[2] C.-H. Chen, B. Karanov, I. Nikoloska, W. van Houtum, Y. Wu, and A. Alvarado, "Modified Baum-Welch Algorithm for Joint Blind Channel Estimation and Turbo Equalization," *International ITG Conference on Systems, Communications and Coding (SCC)*, pp. 1-6, 2025.
[3] B. Karanov, C.-H. Chen, Y. Wu, A. Young, and W. van Houtum, "Data-driven symbol detection for intersymbol interference channels with bursty impulsive noise," *ArXiv preprint arXiv:2405.10814*, 2024.
[4] C.-H. Chen, B. Karanov, W. van Houtum, Y. Wu, and A. Alvarado, "Turbo Receiver Design with Joint Detection and Demapping for Coded Differential BPSK in Bursty Impulsive Noise Channels," *ArXiv preprint arXiv:2412.07911*, 2024.

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

Impact of Environmental Conditions on Fingerprint Image Quality and Recognition Performance

Florens de Wit[1]

1. introduction

Though initially an assumption [1], fingerprints have been shown to be permanent and distinctive from birth to death[2], [3] , unless the dermis is scarred or damaged [4], [5]. Nevertheless, wear, scratches and aging do appear to reduce the degree to which the pattern can be reliably captured [6] which suggests the issue lies with changes in the interaction between finger and sensor, and with captured image quality, instead of changes to the pattern itself. One common impactful factor for this interaction is skin moisture [7], where both excessively dry and wet skin impact the degree to which the fingerprint recognition is still reliable. This directly explains any variation of fingerprint recognition performance where a change or difference in skin dryness is expected (e.g. across age groups [6], [8]) and indirectly explains variation due to environmental conditions through their impact on skin moisture content and surface moisture.

The degree to which moisture related environmental factors impact performance should therefore be of interest to developers, operators and users of biometric systems that use the fingerprint modality. Despite this, most studies on the impact of environmental factors are either focussed on a particular application (e.g. [9]) or focus on a specific context such as certification ([10] further expanded in [11]).

2. Proposed study

We therefore propose an experimental study in which we capture fingerprints of different subjects with the sensor placed in an enclosure kept at different combinations of temperature and humidity.
The conceptual design of the setup we will use is a transparent box that contains a fingerprint scanner, and the means to both independently cool/heat and dry/moisten the air within it. For both temperature variation and drying we will use so called Peltier elements; for moistening we will use an ultrasonic atomizer. Any person which fingerprints are to be scanned can reach into the box through a hole and use the fingerprint scanner. This allows for the acquisition of fingerprints under different environmental conditions without putting subjects in such conditions as a whole.
The preliminary protocol we propose to use is represented as a flow chart in Figure 1. Its main objective is to build a dataset of fingerprints from different people, of multiple fingers of each hand, captured under different combinations of air temperature and relative humidity.

3. Pilot study

To get a first impression of what differences in quality and recognition performance we should expect we performed a pilot experiment that consists of the following:

Using the author as the only subject,

1. Assess quality variation between fingers and/or treatments by;
    a. Scanning thumb, index, middle and ring finger of both hands (10x each finger)
        i. Without treatment
        ii. With skin moisturising cream
        iii. With moisture on skin by breathing on fingers
    b. Inspect images for apparent differences between fingers and/or treatment
    c. Calculating NFIQ2 [12] and visualize (+ observations)

2. Assess variation in verification performance between fingers and/or treatments by
    a. enrolling thumb, index, middle and ring finger from both hands (without treatment)
    b. verification of all fingers (10x each finger)

[1] Data Management & Biometrics, Electrical Engineering, Mathematics & Computer Science, University of Twente

i. without treatment
ii. with skin moisturising cream
iii. with moisture on skin by breathing on fingers

c. document maximum verification score and visualize (+ observations)

3. Use visualisations from 1c and 2c to assess NFIQ vs. recognition

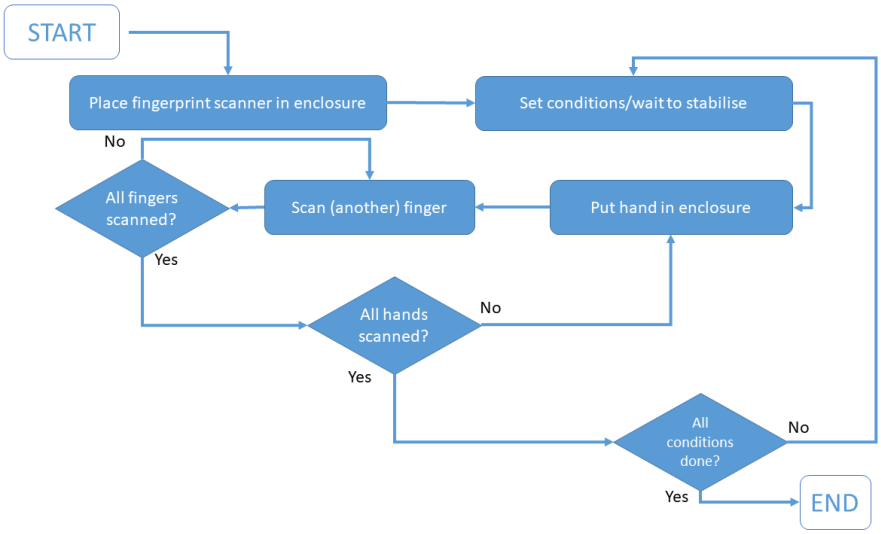

Figure 1: proposed acquisition protocol flow chart

4. observations & results

As can be seen in Figure 2, the amount of moisture added to the fingerprint increases the degree to which the ridges are darkened. This is true both when comparing the same fingerprint (as is here) and for fingerprints in general.

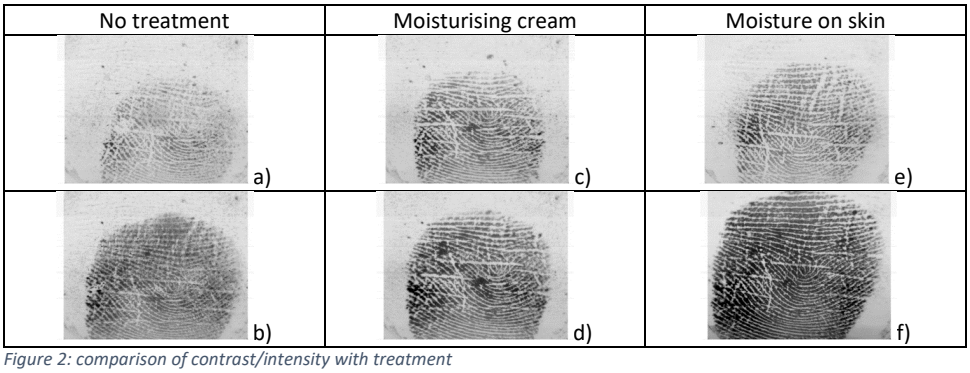

Figure 2: comparison of contrast/intensity with treatment

As can be seen in Figure 3, however, the NFIQ2 quality score does not improve over-all; its distribution appears to broaden and its median moves slightly upward.

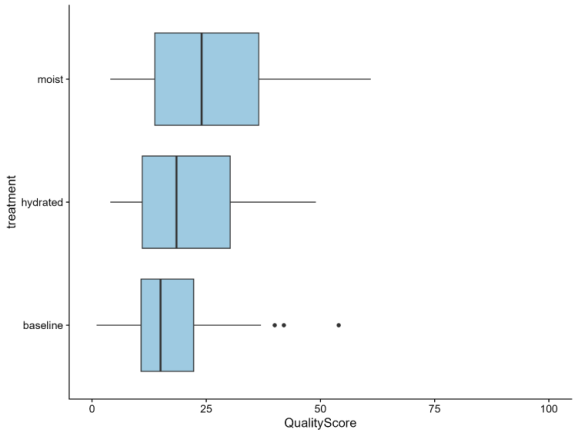

Figure 3: boxplot NFIQ2 quality score by treatment

However on individual fingers the treatments appear to have a very different impact; Figure 4 shows this in the very different behaviour of the NFIQ2 score between treatments. Notice that the hands appear to be offset with respect to each other, with the left hand consistently giving lower quality scores. Even similar fingers are not showing a similar change between treatments, e.g. the right thumb shows a consistent rise in quality score, while the left is almost constant. In fact, all right-hand fingers show a consistent rise in quality with moisture level, while there is a mixed response in left-hand fingers.

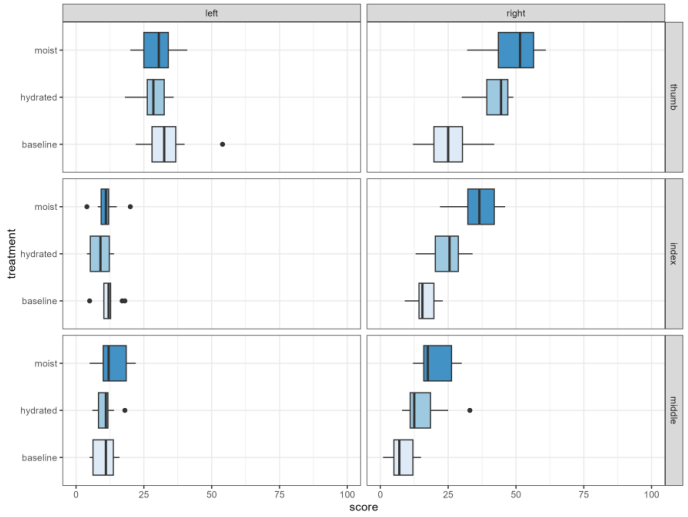

Figure 4: boxplot NFIQ score by treatment for individual fingers

If we consider that the NFIQ is intended as a prediction for the true match rate and compare the verification score by treatment (Figure 5) with the corresponding NFIQ2 score (Figure 3), we see a fairly good correspondence of the quality score with the degree of true matches. If we distinguish between individual fingers (Figure 4/Figure 6), however, the correspondence appears not so good; e.g. the left thumb should have the highest quality but clearly does not.

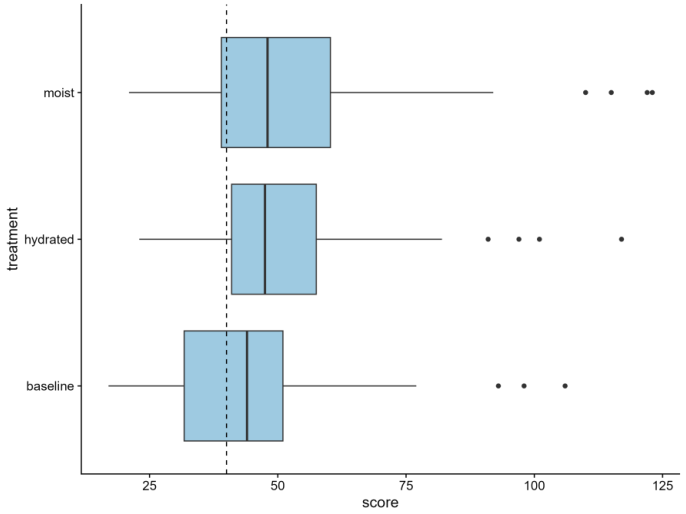

Figure 5: Verification score by treatment; dotted line is threshold for verification algorithm used

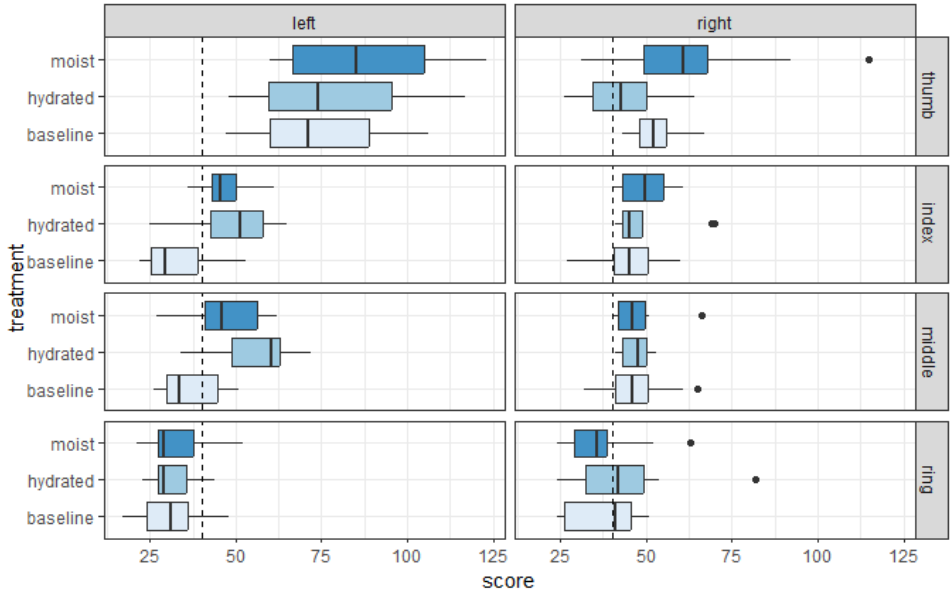

Figure 6: Verification score by treatment for individual fingers; Dotted line is threshold for verification algorithm used

5. discussion & conclusions

Since the results are based on a limited dataset acquired from a single person, we take great caution in drawing any conclusion from them. It is likely there are individual biases present. Especially the variation in results between individual fingers seen in this dataset requires some further investigation, with data from multiple subjects.

We also note that NFIQ 2, just as its predecessor NFIQ is intended to be applied to fingerprint images of a particular resolution and other properties. This may require the images from commonly used scanners to be re-scaled, which may introduce bias.

The verification scores have been calculated directly from live scanned images and live enrolled references, and may therefore not reflect the performance on the same dataset as the NFIQ2 has been calculated. We intend to calculate verification scores from previously scanned images once we start acquiring data using the proposed "climate controlled" setup.

Despite these caveats, the results do indicate we need to critically evaluate whether NFIQ2 is an accurate predictor of verification performance when scanners are used that do not deliver images within the current scope of this quality algorithm.

References

[1] H. Faulds, 'On the Skin-Furrows of the Hand', *Nature*, vol. 22, no. 574, pp. 605–605, Oct. 1880, doi: 10.1038/022605a0.
[2] Didier Meuwly, 'De Vingerafdruk', in *Forensische Wetenschap*, Deventer: Kluwer, 2008, pp. 323–344.
[3] C. Champod, *Fingerprints and other ridge skin impressions*. Boca Raton, FL: CRC Press, 2004.
[4] M. Okajima, 'A methodological approach to the development of epidermal ridges viewed on the dermal surface of fetuses', *Prog. Clin. Biol. Res.*, vol. 84, pp. 175–188, 1982.
[5] Soweon Yoon, Jianjiang Feng, and A. K. Jain, 'Altered Fingerprints: Analysis and Detection', *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 3, pp. 451–464, Mar. 2012, doi: 10.1109/TPAMI.2011.161.
[6] European Commission. Joint Research Centre., *Automatic fingerprint recognition: from children to elderly : ageing and age effects.* LU: Publications Office, 2018. Accessed: Apr. 23, 2025. [Online]. Available: https://data.europa.eu/doi/10.2760/809183
[7] M. A. Olsen, M. Dusio, and C. Busch, 'Fingerprint skin moisture impact on biometric performance', in *3rd International Workshop on Biometrics and Forensics (IWBF 2015)*, Gjovik, Norway: IEEE, Mar. 2015, pp. 1–6. doi: 10.1109/IWBF.2015.7110223.
[8] J. Galbally, R. Haraksim, and L. Beslay, 'A Study of Age and Ageing in Fingerprint Biometrics', *IEEE Trans. Inf. Forensics Secur.*, vol. 14, no. 5, pp. 1351–1365, May 2019, doi: 10.1109/TIFS.2018.2878160.
[9] R. F. Stewart, M. Estevao, and A. Adler, 'Fingerprint recognition performance in rugged outdoors and cold weather conditions', in *2009 IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems*, Washington, DC, USA: IEEE, Sep. 2009, pp. 1–6. doi: 10.1109/BTAS.2009.5339061.
[10] A. Wone, J. D. Manno, C. Charrier, and C. Rosenberger, 'Impact Of Environmental Conditions On Fingerprint Systems Performance', in *2021 18th International Conference on Privacy, Security and Trust (PST)*, Auckland, New Zealand: IEEE, Dec. 2021, pp. 1–5. doi: 10.1109/PST52912.2021.9647754.
[11] A. Wone, 'Contribution to the certification of fingerprint systems: towards the reproducibility of the evaluation', PhD thesis, Normandie Université, Caen, 2023. [Online]. Available: https://theses.hal.science/tel-04431507/
[12] E. Tabassi *et al.*, 'NFIQ 2 NIST fingerprint image quality', National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST IR 8382, Jul. 2021. doi: 10.6028/NIST.IR.8382.

Accelerating Selective Sweep Detection using AMD Deep Learning Processing Units and Vitis AI

Sebastian Bunda*, Nikolaos Alachiotis*, Luuk Spreeuwers*
*University of Twente
{s.t.bunda, n.alachiotis, l.j.spreeuwers}@utwente.nl

***Abstract*—Recent advancements in artificial intelligence have made the need for faster computation through AI acceleration increasingly important. This work explores using AMD's Deep Learning Processor Units and Vitis AI to accelerate an image classification problem to identify a reduction in genetic variation, also known as selective sweep detection. An existing CNN designed for selective sweep detection is investigated and modified for faster image processing. The deployment of the original model showed slow processing due to several layers being handled by the CPU. By modifying the network and removing problematic layers, inference times improved significantly with minimal accuracy loss. The results suggest that DPUs are effective when models avoid custom or unsupported layers. Vitis AI simplifies quantization and compilation, but can be challenging to debug and requires specific software versions, limiting flexibility. Running the software on a cluster without root access also posed difficulties.**

***Index Terms*—Selective Sweep, Convolutional Neural Networks, AI acceleration, Deep Learning Processor Unit, Vitis AI**

I. Introduction

Positive selection is a key driver of genetic diversity, influencing the evolutionary trajectory of populations. When a beneficial mutation increases in frequency and therefore reduces the genetic variation in the surrounding region, this is known as a selective sweep. This serves as a key indicator for a positive selection. Mutations in the genome within a population can be visualized in a 2D image representation where a black pixel is a mutation and white pixels are no change to a baseline genome, where each row in the image is another individual in the population. Since this data can be represented in an image representation, a selective sweep can be classified using a Convolutional Neural Network, as in SweepNet [1], Fast-NN [2], and FASTER-NN [3].

Recent developments in artificial intelligence have demanded significant improvements in computational efficiency and performance. A key driver in this is the use of AI accelerators [4], [5], specialized hardware to speed up key operations within Deep Neural Networks (DNNs) such as matrix multiplications and tensor operations.

To compete with the popular GPU AI acceleration trend, AMD tried to join the market of AI Accelerators in the form of AI Engines, to promote very efficient and high-throughput AI inference. The Versal Adaptive System on Chip platform combines programmable logic and AI Engines to create a new type of FPGA specialized in AI acceleration. To improve the accessibility to non-hardware engineers, AMD presented Deep Learning Processor Units (DPU), which are programmable engines dedicated to convolutional networks that can run on AMD FPGAs and are designed to take advantage of the AI-Engines when available [6].

In this work, we have investigated the use case for DPU to accelerate selective sweep detection using SweepNet [1]. We implemented the DPU on both the Versal VCK5000 card and the KV260 FPGA using Vitis AI.

II. Background

A. Vitis AI Framework

Vitis AI [7] is an environment for rapid prototyping and implementing AI inference on AMD's FPGA and adaptive SoC devices. The primary objective of this software and the framework is to manage and simplify the complexities of hardware implementation and minimize the need for hardware-specific knowledge. The framework features an optimizer capable of network pruning, albeit with occasional limitations in functionality. Furthermore, Vitis AI includes a quantizer that facilitates the conversion of standard 32-bit floating-point models to 8-bit integer models. Additionally, it is equipped with a compiler that translates the quantized model description into a DPU instruction sequence. Finally, the framework includes a comprehensive model zoo, which includes pretrained and compiled models such as MobileNet and ResNet for different DPUs.

B. Deep Learning Processor Units

The Deep Learning Processor Unit (DPU) [6] is a soft-core accelerator overlay that accesses the programmable logic (PL) of the FPGA or the AMD AI Engines when available on the hardware. The DPU has a specialized instruction set, for which the Vitis AI Compiler compiles an *.xmodel* from the original neural network in C or Python. Each DPU can have different instructions on how to process a DNN layer based on the resources available on the hardware device.

III. Method

The basis of this work is the CNN architecture SweepNet[1] used to classify images as either Neutral or Selection. The SweepNet architecture can be seen in Fig. 1.

[1] https://github.com/Zhaohq96/SweepNet

A. DPUCVDX8H

The DPU used for the VCK5000 card, a device with AMD's dedicated AI Engines, was the DPUCVDX8H-8PE. The focus of this DPU is high-throughput inference and can process 8 images in parallel. For the VCK5000 card, there were four versions available: the 4PE, 6PE-DWC, 6PE-MISC and the 8PE, as shown in Table I. The main difference was that the 4PE processed four images in parallel, the 6PE processed six images, and the 8PE processed eight images. In our experience, there was not much of a processing time difference per image for each of the DPUs for our neural network. It was decided to use the DPU with the most processing power.

TABLE I
THE RESOURCES FOR EACH DPU OPTION FOR THE DPUCVDX8H [8] AND THE DPUCZDX8G [9]

DPU	Arch	LUT	Register	Block RAM	URAM	DSP
DPUCVDX8H	4pe_miscdwc	221K	254K	456	312	44
	6pe_misc	223K	287K	684	390	34
	6pe_dwc	631K	648K	780	402	424
	8pe_normal	**674K**	**696K**	**912**	**424**	**524**
DPUCZDX8G	**B4096**	**52K**	**99K**	**0**	**68**	**710**

B. DPUCZDX8G

The processing core of the Kria KV260, the DPUCZDX8G, is highly configurable. However, we found that not all configurations were able to be built successfully using Vitis. For the project, we generated the *DPUCZDX8G_ISA1_B4096_0101000047010407* without problems. The specifications compared to the DPUCVDX8H-8PE can be found in Table I.

C. Quantization and Compilation

To run our network on the two devices, the network had to be quantized to an 8-bit integer representation. Since the original network was in a 32-bit representation, this can lead to some small reduction in accuracy. The quantizer used in this project was Post-Training Quantization, which means that the network is quantized after training and is calibrated using a set of unlabeled data.

After the network is quantized, it can be compiled for the specific DPU. This means that the computational graph, together with the control and data flow, will be constructed exploiting the available hardware for optimal parallelism and data re-use.

D. Deploying the Original Model

Using the Vitis AI workflow [10], the SweepNet network [1] can be quantized, compiled and deployed on the VCK5000 card and the KV260 FPGA. This resulted in an image per second of 9.9 and 67.7, respectively. The DPU supports several types of layers that can be found in convolutional neural networks. Unfortunately, the SweepNet architecture has a Multiplication and Global Average Pooling layer, which are not supported by the DPUCVDX8H. The DPU does support Average Pooling, but with a maximum kernel size of 7. Within the SweepNet architecture, the feature maps of 122x122 need to be averaged, which is too large for the DPU. The compilation logs showed that the Global Average Pooling and the Multiply layers, in particular, in the original network, were allocated to the CPU on the devices. This meant that several subgraphs were created and thus slowing down the processing considerably. These layers were not supported by the DPUCVDX8H and DPUCZDX8G DPUs.

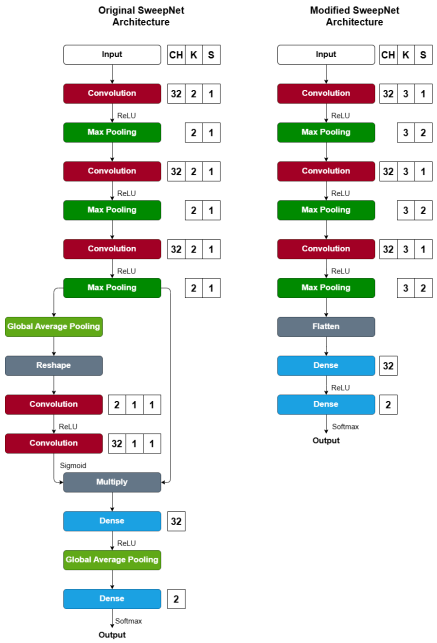

Fig. 1. The SweepNet architecture and the Modified SweepNet architecture. The padding also changed from 'valid' to 'same'. CH, K and S mean the number of output channels, kernel size and stride respectively.

E. Modified SweepNet

To optimize the SweepNet architecture for the two devices without compromising accuracy, several modifications were made. The original network included a Squeeze and Excitation (SE) block with a Global Average Pooling layer and a multiplication layer. Due to the DPU's limitation of supporting average pooling layers with a maximum kernel size of 7, and the GAP layer requiring a kernel size of 122, the SE block was omitted, which initially led to poor performance due to a reduced receptive field. To address this, the kernel size of the preceding convolutional layers was increased from 2 to 3, which seemed to retain the information better.

Furthermore, we decided to reduce the feature maps by adding a stride to the pooling layers instead of using no padding. Lastly, the final Global Average Pooling layer was removed by flattening the output of the convolutional layers before the first dense layer. The modifications did increase the model size considerably (281K vs. 9.7K parameters).

TABLE II
THE ACCURACIES FOR THE ORIGINAL SWEEPNET ARCHITECTURE AND OUR MODIFIED ARCHITECTURE. THE D1 AND D2 DATASETS ARE DESCRIBED IN THE ORIGINAL WORK [1]. FP32 IS THE ORIGINAL MODEL, AND INT8 IS THE QUANTIZED MODEL.

	D1		D2	
Architecture	Accuracy FP32	Accuracy INT8	Accuracy FP32	Accuracy INT8
SweepNet [1]	99.70%	99.70%	93.6%	93.55%
Modified SweepNet (ours)	99.85%	99.65%	93.7%	94.05%

IV. EXPERIMENTAL DESIGN

The modified network will be assessed for its accuracy in detecting selective sweeps using both the 32-bit trained model and the quantized 8-bit model. To assess the acceleration using the AI-Accelerators, the inference speed will be compared with a commercial CPU and GPU.

A. Quantization Performance

To assess the selective sweep classification performance of the Modified SweepNet architecture, the same data is used as in the original paper [1]. It was decided not to do the pre-processing steps to focus solely on the performance of the neural network, and it didn't seem to affect the accuracy of the original network significantly. The results can be found in Table II. The two datasets used in this experiment are based on a simulated population bottleneck, which results in a reduced genetic diversity [3]. The D1 dataset has a mild population bottleneck and the D2 dataset has a severe population bottleneck and is thus a more difficult classification task.

B. Acceleration Performance

The acceleration performance is assessed by comparing the inference processing time for a single image on four devices: a consumer CPU (AMD Ryzen 7 Pro 8840U), a consumer GPU (Nvidia GTX 1060), the KV260 and the VCK5000. The performance for the CPU and GPU will be compared for different batch sizes, where the inference time will be computed by multiplying the time to process a batch by the batch size.

The Kria KV260 was only able to process one image at a time, and for the VCK5000, we found that the 8PE DPU was the fastest. The results can be found in Fig. 2.

V. RESULTS ANALYSIS

Our experiments show that when the global average pooling and multiply layers were removed and some slight modifications were made to compensate, as seen in Fig. 2 and Table II, the inference time can be significantly reduced with a minimal reduction in accuracy (in the D2 case, it was even slightly better). When the modified network was quantized and compiled for the VCK5000 and the KV260, the image per second was now 4009 (with 8 images at the same time) and 1615 (for 1 image at a time), respectively. The results indicate that when models do not incorporate custom or unsupported layers, the DPU can be utilized effectively (as only one computation graph was compiled).

The VCK5000 seems to process the images in batches of 8 the fastest, compared to the CPU and the GPU. It is interesting to see how the GPU does not seem to improve with bigger batch sizes than 8. The memory of 3GB is likely the limiting factor here, and with higher memory bandwidth and more memory, newer GPUs are likely to be able to outperform the VCK5000 without further optimizations.

Table I shows a much larger available compute resource for the DPUCVDX8H compared to the DPUCZDX8G; however, the results from Fig. 2 don't show a much larger improvement in speed. This suggests that the hardware of the VCK5000 is capable of much larger neural network layers, such as those required by Large Language Models or Foundational Models. This also indicates that the compute capabilities of the DPUs are fully utilized, so there might be more room for improvement.

Another interesting observation that can be made from Fig. 2 is the fact that even though the modified architecture is about 30 times larger, it is still faster on all devices compared to the original architecture. We expect that the multiplication layer in the original model is not computed efficiently on all hardware platforms.

A remarkable observation is also the fact that the KV260 FPGA has the highest single-image processing speed. In use cases where only one image has to be processed at the same time, the KV260 outperforms all other devices. It is also interesting to note that the KV260 only has a 36W power rating [11], making it very suitable for embedded computer vision applications.

In future work, more research needs to be done on which specific layers impact the performance for each DPU, to be able to design a model tailored to the DPU specifically. In parallel, we can continue the work of FASTER-NN [3] to exploit the information within the data specifically and more carefully design the neural architecture for the application.

A. Vitis AI

The black-box approach of Vitis AI simplifies the processes of quantization and compilation, making them straightforward but challenging to debug. Other studies [12], [13] have demonstrated superior inference times in other applications when using custom accelerators, which necessitate a deeper understanding of the hardware. The Vitis AI method, however, largely circumvents the need for such specialized knowledge. Additionally, it is important to note that the devices used in this experiment required particular versions of Vitis AI, which can limit flexibility. To get the software working on a cluster without root access was also problematic and required considerable effort.

VI. CONCLUSION

The goal of this research was to explore AMD's Deep Learning Processor Units in the context of accelerating a convolutional neural network implementation for selective sweep classification. The experiments show that the DPUs can

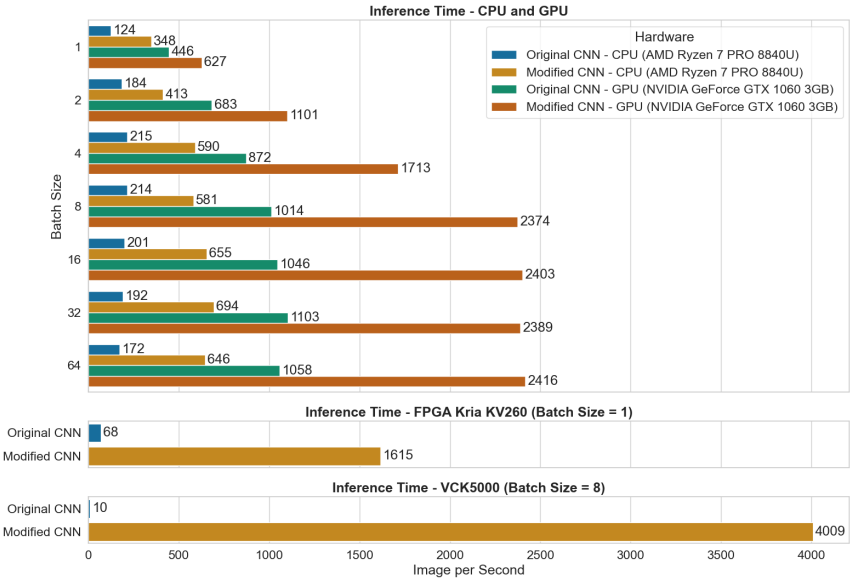

Fig. 2. The inference time for a single image for the original architecture and the modified architecture. The inference time is the average over 5 iterations on a CPU, GPU, the Kria KV260 and the VCK5000. The Image per Second was calculated by multiplying the inference time of one batch by the batch size.

be very effective as an acceleration platform, on the condition that all the layers are supported and within the specification limits. The most remarkable result is that the Kria KV260 had the fastest single-image processing performance, making the FPGA a very attractive platform for embedded AI solutions.

The Vitis AI framework is a very effective tool that enables the use of the AMD devices as AI accelerators without requiring expert knowledge of reconfigurable hardware. However, the software does require the user to limit the design options considerably, and when the software does not work, it is very difficult to debug due to the black-box approach, which then requires knowledge of the hardware.

ACKNOWLEDGMENTS

I want to give a special thanks to Dorus Abeln for the technical support and to Sjoerd van den Belt and Hanqing Zhao for the helpful discussions.

REFERENCES

[1] H. Zhao, P. Pavlidis, and N. Alachiotis, "Sweepnet: A lightweight cnn architecture for the classification of adaptive genomic regions," in *Proceedings of the Platform for Advanced Scientific Computing Conference*, ser. PASC '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3592979.3593411

[2] S. van den Belt, H. Zhao, and N. Alachiotis, "Scalable CNN-based classification of selective sweeps using derived allele frequencies," vol. 40, pp. ii29–ii36. [Online]. Available: https://doi.org/10.1093/bioinformatics/btae385

[3] S. van den Belt and N. Alachiotis, "Fast and accurate deep learning scans for signatures of natural selection in genomes using FASTER-NN," vol. 8, no. 1, pp. 1–12, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s42003-025-07480-7

[4] A. Reuther, P. Michaleas, M. Jones, V. Gadepally, S. Samsi, and J. Kepner, "Ai and ml accelerator survey and trends," in *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, 2022, pp. 1–10.

[5] G. Akkad, A. Mansour, and E. Inaty, "Embedded deep learning accelerators: A survey on recent advances," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 5, pp. 1954–1972, 2024.

[6] "DPU for Convolutional Neural Network." [Online]. Available: https://www.xilinx.com/products/intellectual-property/dpu.html

[7] *Vitis AI Developer Hub*, Advanced Micro Devices, 2025, available at https://www.amd.com/en/developer/resources/vitis-ai.html.

[8] *DPUCVDX8H for Convolutional Neural Networks Product Guide (PG403)*, Advanced Micro Devices, January 2023, available at https://docs.amd.com/r/en-US/pg403-dpucvdx8h.

[9] *DPUCZDX8G for Zynq UltraScale+ MPSoCs Product Guide (PG338)*, Advanced Micro Devices, January 2023, available at https://docs.amd.com/r/en-US/pg338-dpu.

[10] *Vitis AI 3.0 User Guide (UG1414)*, Advanced Micro Devices, Febuary 2023, available at https://docs.amd.com/r/3.0-English/ug1414-vitis-ai.

[11] *Kria KV260 Vision AI Starter Kit Data Sheet*, Advanced Micro Devices, 2022, available at https://docs.amd.com/r/en-US/ds986-kv260-starter-kit/Product-Details.

[12] M. Machura, M. Danilowicz, and T. Kryjak, "Embedded object detection with custom littlenet, finn and vitis ai dcnn accelerators," *Journal of Low Power Electronics and Applications*, vol. 12, no. 2, 2022. [Online]. Available: https://www.mdpi.com/2079-9268/12/2/30

[13] D. K. Athur, R. Pawar, and A. Arora, "Out-of-the-box performance of fpgas for ml workloads using vitis ai," in *Applied Reconfigurable Computing. Architectures, Tools, and Applications: 21st International Symposium, ARC 2025, Seville, Spain, April 9–11, 2025, Proceedings*. Berlin, Heidelberg: Springer-Verlag, 2025, p. 123–139. [Online]. Available: https://doi.org/10.1007/978-3-031-87995-1_8

Transformer-based Motor Fault Detection Using DWT

Jiarui Zhou [1], Sinian Li [1], Edmund Marth [2], Patrick Zorn [2], Wolfgang Gruber [2], Justin Dauwels [1]
[1] *Signal Processing Systems, Dept. of Microelectronics, TU Delft, 2628 CD, The Netherlands*
[2] *Institute of Electric Drives and Power Electronics, Johannes Kepler University Linz, Linz, Austria*

Abstract

Recent advances in mechanics and autonomous control have made industrial manufacturing systems increasingly complex and thus more prone to faults, degradation, and unexpected breakdowns. Early detection of faults is crucial to reduce downtime, perform timely maintenance, and increase productivity. Recent advances in powerful feature extraction techniques and machine learning models have enabled unmanned, intelligent fault detection systems [1]. Motivated by the demand for highly accurate and robust fault detection, we propose a Transformer-based model utilizing Discrete Wavelet Transform (DWT) for feature extraction and validate its performance using a test-bench dataset.

In this paper, the fault detection task is divided into two steps: feature extraction using DWT decomposition and Transformer-based classification, as shown in Figure 1. In the first step, the signal is decomposed into wavelet coefficients using DWT. The key advantage over Fourier transform is its flexible temporal resolution. It can capture features in both time domain and frequency domain. DWT applies high-pass and low-pass filters to extract detail and approximation coefficients, respectively, followed by down-sampling to remove redundancy. This multi-level decomposition enhances frequency resolution, preserving low-frequency components and isolating high-frequency details. The resulting coefficients can be used as features for machine learning models effectively. In the second step, a Transformer-based model performs classification using attention mechanism to capture key parts of long sequences. The wavelet coefficients from the previous step, which are long time series of varying lengths, are individually processed by separate encoders. A learnable class token is prepended to each sequence as a representation. After encoding, these learnable class tokens are concatenated to form a new feature, which is converted into a class prediction through a multi-layer perceptron. The dataset was collected using a test bench consisting of a motor, a torque sensor, a hysteresis brake, and power electronics. Misalignment of the hall sensors that can lead to a commutation angle error serves as the failure class. Three-phase current signals are the input to our model.

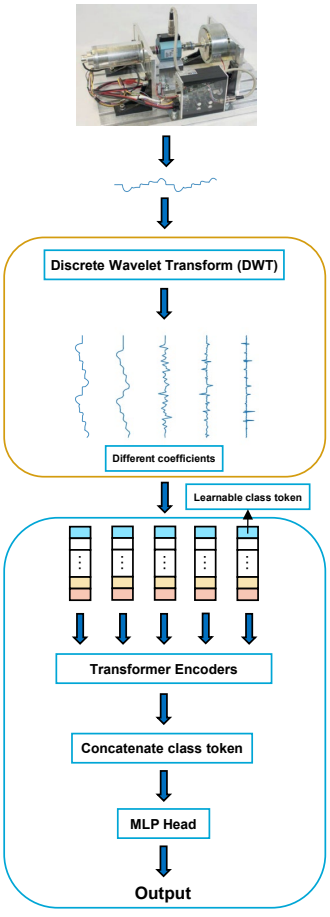

Figure 1. Pipeline of the proposed model

The proposed model is compared with convolutional neural networks (CNN), commonly used in classification tasks and long short-term memory (LSTM) networks that are basic for time series analysis. The performance comparison is as follows: The GRU model obtained an accuracy of 94.25%, recall of 92.88%, and precision of 95.30%. The 1D-CNN and LSTM models both achieved an accuracy of 96.75%, recall of 93.89%, and precision of 99.46%. The proposed model outperforms all others, achieving an accuracy of 97.25%, recall of 94.40%, and precision of 100%. This demonstrates the effectiveness of the model on the target dataset. The integration of DWT as a feature extractor facilitates the extraction of multi-scale signal information. The Transformer is able to learn features from long time series, helping the model effectively capture and classify low- and high-frequency signal characteristics. In future research, we plan to explore the generalization capability of the model, given that only one dataset was considered in current study.

References

[1] Saufi, S. R., et al., "Challenges and opportunities of deep learning models for machinery fault detection and diagnosis: A review," *Ieee Access*, Vol. 7, 2019, pp. 122644–122662.

This work is part of R-PODID project, supported by the Chips Joint Undertaking and its members, including the top-up funding by National Authorities of Italy, Turkey, Portugal, The Netherlands, Czech Republic, Latvia, Greece, and Romania under grant agreement n° 101112338.

Revisiting Rotational Cryptanalysis on chained modular additions

Addie Neyt

Rotational cryptanalysis, a chosen-plaintext attack introduced by Khovratovich and Nikolić [KN10], analyzes how rotational pairs $(x, S^r(x))$, where $S^r(x)$ denotes a rotation by a fixed constant r, propagate through ARX (Addition modulo 2^n, Rotation, XOR) operations. Their initial work demonstrated that both XOR and rotation operations preserve rotational pairs with probability one. They further investigated how modular addition affects these pairs, showing that the preservation probability depends solely on the rotation amount. However, their approach faced limitations. The original analysis treated modular additions as independent events and calculated, for a cipher containing q additions each preserving a rotational pair with probability p_r, the overall probability of preserving the rotational input as p_r^q.

This assumption breaks down in the presence of chained modular additions (the output of one modular addition is chained as one of the inputs for the next modular addition), which introduces dependencies that influence the actual preservation probability. This issue was addressed in a follow-up study [KNP+15], where they showed that chained modular additions in ARX ciphers do not form a Markov chain in the context of rotational analysis. Consequently, the rotational probability cannot be computed as a simple product of individual probabilities. Their revised analysis focused on the case where all outputs of the chained additions are required to be rotational, which they state to be an important requirement as in ARX, outputs of intermediate modular additions are used as inputs to other operations and are assumed to be rotational. They also noted that if only the final output of a chain of additions needs to maintain a rotational pair, a different formula applies, though it was omitted due to space constraints.

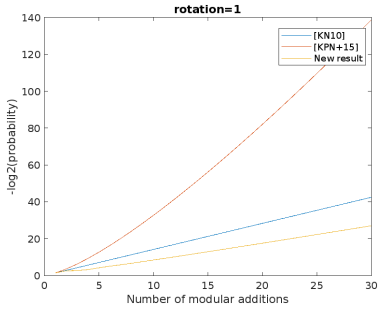

However, this requirement is not the case for all current ARX ciphers. For example, in SHA-2, one can identify chains of modular additions where the intermediate results are not used anywhere else in the cipher. In such cases, it suffices that only the final output of the chain preserves the rotational pair. We derived the rotational probability for this relaxed setting and found that it is significantly higher compared to the stricter assumption where all intermediate values must be rotational.

Although we did not conduct a detailed analysis of rotational cryptanalysis on SHA-2 using this adjusted model, preliminary results suggest that the rotational probabilities differ substantially when using this new formula. However these findings do not yet rival the best-known attacks on SHA-2. They do highlight an important consideration: cipher designers must be mindful of whether intermediate values between modular additions are used as input for other operations, as it greatly influences the effectiveness of rotational cryptanalysis.

References

[KN10] Dmitry Khovratovich and Ivica Nikolic. Rotational cryptanalysis of ARX. In Seokhie Hong and Tetsu Iwata, editors, *FSE 2010*, volume 6147 of *LNCS*, pages 333–346. Springer, Berlin, Heidelberg, February 2010.

[KNP+15] Dmitry Khovratovich, Ivica Nikolic, Josef Pieprzyk, Przemyslaw Sokolowski, and Ron Steinfeld. Rotational cryptanalysis of ARX revisited. In Gregor Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 519–536. Springer, Berlin, Heidelberg, March 2015.

A Appendix

Notations

- $S^r(x)$ is the rotation by the fixed constant r
- $\boxplus$ is the modular addition modulo 2^n

Theorem 1. *Let $a_1, \ldots, a_k$ be n-bit words chosen at random and let r be a positive integer such that $0 < r < n$. Then*

$$\begin{aligned}\Pr[S^r(a_1 \boxplus a_2 \boxplus \ldots \boxplus a_k)] &= \Pr[S^r(a_1) \boxplus S^r(a_2) \boxplus \ldots \boxplus S^r(a_k)] \\ &= P_{C_r} \cdot P_{C_{n-r}}\end{aligned}$$

We define

1. *$a_i = x_i || y_i$, with x_i an r-bit word, y_i an $(n-r)$-bit word and $||$ concatenation*
2. *C_r is the carry of the sum $x_1 \boxplus x_2 \boxplus \ldots \boxplus x_k$ and $P_{C_r} = \Pr[C_r \equiv 0 \mod 2^{n-r}]$*
3. *C_{n-r} is the carry of the sum $y_1 \boxplus y_2 \boxplus \ldots \boxplus y_k$ and $P_{C_{n-r}} = \Pr[C_{n-r} \equiv 0 \mod 2^r]$*

$$P_{C_r} = \begin{cases} 2^{-rk}\binom{k+2^r-1}{2^r-1} & \text{if } 2^{n-r} > k-1 \\ 2^{-rk}\binom{k+2^r-1}{2^r-1} + \sum\limits_{t'=2^{n-r},\ t' \text{ multiple of } 2^{n-r}}^{\lfloor \frac{k-1}{2^{n-r}} \rfloor 2^{n-r}} f(t', 2^r) & \text{else} \end{cases}$$

$$P_{C_{n-r}} = \begin{cases} 2^{-(n-r)k}\binom{k+2^{n-r}-1}{2^{n-r}-1} & \text{if } 2^r > k-1 \\ 2^{-(n-r)k}\binom{k+2^{n-r}-1}{2^{n-r}-1} + \sum\limits_{t'=2^r,\ t' \text{ multiple of } 2^r}^{\lfloor \frac{k-1}{2^r} \rfloor 2^r} f(t', 2^{n-r}) & \text{else} \end{cases}$$

$$f(t,l) = 2^{-k} l^{-1} \left[\sum_{j=tl}^{\min((t+1)l,\ k(l-1))} \left(\binom{j+k-1}{k-1} - \sum_{m=1}^{\lfloor \frac{j}{l-1} \rfloor} (-1)^{m+1} \binom{k}{m} \binom{j-ml+k-1}{k-1} \right) \right]$$

Proof. We rewrite:

$$S^r(a_1 \boxplus a_2 \boxplus \ldots \boxplus a_k) = S^r(a_1) \boxplus S^r(a_2) \boxplus \ldots \boxplus S^r(a_k) \tag{1}$$

The left side is equal to:

$$\begin{aligned} S^r(a_1 \boxplus a_2 \boxplus \ldots \boxplus a_k)) &= S^r(x_1 \boxplus \ldots \boxplus x_k \boxplus C_{n-r}) || (y_1 \boxplus \ldots \boxplus y_k)) \\ &= (y_1 \boxplus \ldots \boxplus y_k) || (x_1 \boxplus \ldots \boxplus x_k \boxplus C_{n-r}) \end{aligned}$$

The right side is equal to:

$$\begin{aligned} S^r(a_1) \boxplus S^r a_2) \boxplus \ldots \boxplus S^r(a_k) &= S^r(x_1||y_1) \boxplus S^r x_2||y_2) \boxplus \ldots \boxplus S^r x_k||y_k) \\ &= y_1||x_1 \boxplus y_2||x_2 \boxplus \ldots \boxplus y_k||x_k \\ &= (y_1 \boxplus \ldots \boxplus y_k \boxplus C_r)||(x_1 \boxplus \ldots \boxplus x_k) \end{aligned}$$

We get that (1) $\Leftrightarrow \begin{cases} C_{n-r} \equiv 0 \mod 2^r \\ C_r \equiv 0 \mod 2^{n-r} \end{cases}$

Because of independence, we get

$$\Pr[S^r(a_1 \boxplus a_2 \boxplus \ldots \boxplus a_k)] = \Pr[C_r \equiv 0 \mod 2^{n-r}] \Pr[(C_{n-r} \equiv 0 \mod 2^r] = P_{C_r} \cdot P_{C_{n-r}}$$

[KNP+15, Lemma 2] calculates two probabilities $\Pr[C_r = 0] = 2^{-rk} \binom{k+2^r-1}{2^r-1}$ and $\Pr[C_{n-r} = 0] = 2^{-(n-r)k} \binom{k+2^{n-r}-1}{2^{n-r}-1}$

Since $C_{n-r}, C_r \leq k-1$, we get that the probabilities $\Pr[C_r \equiv 0 \mod 2^{n-r}]$ and $\Pr[(C_{n-r} \equiv 0 \mod 2^r]$ will differ from $\Pr[C_r = 0]$ and $\Pr[C_{n-r} = 0]$ only if $k-1 > 2^{n-r} \vee 2^r$

We calculate $P_{C_{n-r}}$ when $k-1 > 2^r$, the calculation for P_{C_r} when $k-1 > 2^{n-r}$ is similar.

$$\begin{aligned} P_{C_{n-r}} &= P(C_{n-r} \equiv 0 \mod 2^r) \\ &= \Pr[C_{n-r} = 0 \vee C_{n-r} = 2^r \vee \ldots \vee C_{n-r} = t \cdot 2^r] \text{ with } t = \left\lfloor \frac{k-1}{2^r} \right\rfloor \\ &= \Pr[C_{n-r} = 0] + \Pr[C_{n-r} = 2^r] + \ldots \Pr[C_{n-r} = t \cdot 2^r] \text{ (mutually exclusive)} \\ &= 2^{-(n-r)k} \binom{k+2^{n-r}-1}{2^{n-r}-1} + \sum_{t'=2^r,\ t' \text{ multiple of } 2^r}^{\lfloor \frac{k-1}{2^r} \rfloor 2^r} \Pr[C_{n-r} = t'] \end{aligned}$$

with

$$\begin{aligned} &\Pr[C_{n-r} = t] \\ &= \Pr[t \cdot 2^{n-r} \leq y_1 + \ldots + y_k < (t+1) \cdot 2^{n-r}] \\ &= 2^{-(n-r)k} \sum_{j=t \cdot 2^{n-r}}^{\min((t+1)\cdot 2^{n-r}, k\cdot(2^{n-r}-1))} \#\{y_1 + \ldots + y_k = j \wedge 0 \leq y_i < 2^{n-r}\} \qquad (2) \\ &= 2^{-(n-r)k} \left[\sum_{j=t2^{n-r}}^{\substack{\min((t+1)2^{n-r}-1, \\ k(2^{n-r}-1))}} \left(\binom{j+k-1}{k-1} - \sum_{m=1}^{\lfloor \frac{j}{2^{n-r}-1} \rfloor} (-1)^{m+1} \binom{k}{m} \binom{j-m2^{n-r}+k-1}{k-1} \right) \right] \qquad (3) \end{aligned}$$

In step (2) $\min((t+1) \cdot 2^{n-r} - 1, k \cdot (2^{n-r} - 1))$ is used in the sum, because of the maximum value that $y_1 + \ldots + y_k$ can take is limited by the smallest of those two. In step (3), we have that $\binom{j+k-1}{k-1}$ counts all possibilities for $\{y_1 + \ldots + y_k = j\}$ without taking into account $0 \leq y_i < 2^{n-r}$. So we have to subtract $\#\{A_1 \cup A_2 \cup A_3 \cup \ldots\}$, with A_i all possibilities for $\{y_1 + \ldots + y_k = j\}$ and i y's that are $\geq 2^{n-r}$. To find this we use the inclusion-exclusion principle of combinatorics.

□

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

Learnable Model Compression for Edge Inference

Joris van de Weg, Sinian Li, Justin Dauwels
Signal Processing Systems, Dept. of Microelectronics, TU Delft, The Netherlands

Abstract

With the growing developments in Artificial Intelligence (AI), the deployment of Deep Learning (DL) models has become an increasingly attractive solution for industrial applications, such as machine health monitoring and predictive maintenance. To achieve real-time analysis and reduce the dependency on cloud infrastructure, it is often more practical and desirable to put the processing of the machine's sensor data directly on edge devices. However, these devices often have limited memory and computational resources. Shallow machine learning models with expert features are small in terms of computations and memory, and thus well-suited for deployment on edge devices. However, DL models have a greater capacity to learn complex and meaningful patterns directly from the data. Yet, this improved performance comes at the cost of significantly higher memory and computational demands, which can easily exceed the limits of an edge device. So, there is a need for compression of these models, which is a big research area in machine learning, not only for edge deployment. However, little work has focused on the joint compression of model architectures that combine convolution and attention layers, especially for edge deployment. To this end, this work proposes a model compression method combining gated low-rank decomposition and quantization, aimed at more effective model inference on edge devices.

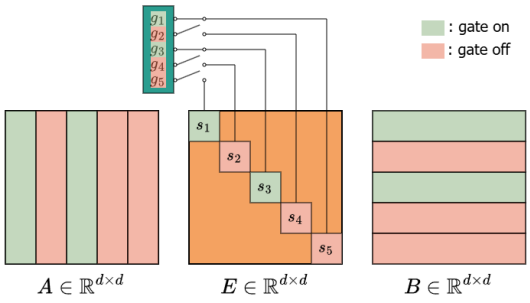

Figure 1. Gated low-rank decomposition, where the matrix is factorized as $A \cdot E \cdot B$ with gating variables controlling the diagonal scaling entries and their associated rows and columns. These gating variables are learned during training to adaptively find the effective rank.

Quantization and low-rank decomposition are among the most effective compression techniques. Counter-intuitively, such compression can even improve generalization, leading to more accurate and robust models. When combined, these techniques can rival or even outperform the baseline models. Despite their benefits, traditional compression methods rely on fixed or manually selected parameters, such as rank thresholds or quantization levels, which can differ for every layer in the model. To address this limitation, our proposed model strategy learns these parameters during training. We follow the Bayesian Bits method [1], which allows the quantization to be adjusted per layer. It makes the quantization controllable during training by rewriting the weight as a summation of different precision levels. For low-rank decomposition, we will extend the LoRA framework by applying it to different decomposition techniques [2]. It uses gating variables to adaptively control the rank by modulating the diagonal components of the decomposition, as shown in Figure 1. Our method is applied to a model architecture containing convolutional layers and attention blocks, as used in the original transformer-based models. Our approach will be evaluated in terms of memory footprint and computational cost, measured in the number of FLOPs. Initial experiments have shown that adaptable quantization can reduce model size by a factor of 12, with accuracy being preserved or enhanced by up to 1 percentage point. While it causes the FLOPs to roughly double, due to the computational overhead of dynamic quantization. Applying low-rank decomposition has shown that it can reduce the memory needs by a factor of 4 while maintaining accuracy, and reducing computations by around 20%. [1]

References

[1] van Baalen, M., Louizos, C., Nagel, M., Amjad, R. A., Wang, Y., Blankevoort, T., and Welling, M., "Bayesian Bits: Unifying Quantization and Pruning," , 2020. URL `https://arxiv.org/abs/2005.07093`.

[2] Zhang, Q., Chen, M., Bukharin, A., He, P., Cheng, Y., Chen, W., and Zhao, T., "Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning," *The Eleventh International Conference on Learning Representations*, 2023. URL `https://openreview.net/forum?id=lq62uWRJjiY`.

[1]This work is part of R-PODID project, supported by the Chips Joint Undertaking and its members, including the top-up funding by National Authorities of Italy, Turkey, Portugal, The Netherlands, Czech Republic, Latvia, Greece, and Romania under grant agreement n° 101112338.

A Sparse Network Design for Fast Multiagent Formation Stabilization

Zhonggang Li, Geert Leus, and Raj Thilak Rajan
Signal Processing Systems, Faculty of EEMCS, Delft University of Technology
Delft, The Netherlands
{z.li-22, g.j.t.leus, r.t.rajan}@tudelft.nl

***Abstract*—Formation control of multiagent systems is an essential task of robotic swarms. Distributed algorithms, such as affine formation control (AFC) [2], typically model the system as a graph in which the vertices and edges denote the agents and their interactions, respectively. For AFC, the stabilization of the system is realized through a consensus-like algorithm using the stress matrix, a generalized graph Laplacian. The stabilizability is guaranteed by the (universal) rigidity of the graph, which is then translated to the positive semi-definiteness (PSD) and rank of the stress matrix. On the other hand, the speed of convergence is related to the smallest nonzero eigenvalue of the stress matrix. For many practical reasons, fewer edges are preferred in the graph design since it helps reduce the communication load but might jeopardize the convergence speed and rigidity, thus making an optimized solution necessary. In this work, we address a network design problem for the AFC framework that optimizes for a proper stress matrix featuring: 1) PSD with the required rank to satisfy the graph rigidity; 2) sparsity such that the number of edges are minimized; 3) a maximized smallest nonzero eigenvalue to promote the speed of convergence.**

Conventionally, universally rigid graph design methods involve either a hand-design of a graph followed by a semidefinite program for the stress matrix [3] or a complex mixed-integer programming that designs the stress matrix directly but has no optimality and computational advantages [4]. In this work, we propose a convex optimization framework where we seek a sparse stress vector from a complete graph using an L1 minimization. We also translate the eigenvalue maximization problem into a set of convex objectives and constraints. We offer some insights and guidance into the choice of the hyperparameters involved in the optimization, and we give a few numerical examples to show that our proposed method can reach a sparse solution with an improved speed of convergence. One of the examples is shown in Fig. 1, where our results are compared with several state-of-the-art methods given a circular configuration. Fig. 2 presents the actual error convergence in the formation control algorithm associated with the smallest eigenvalue in the stress matrix. As a conclusion, our proposed solution can give a sparser network with faster convergence.

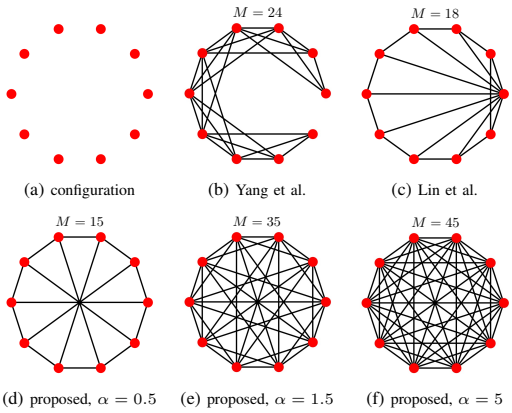

Fig. 1. Resulting graphs of our proposed solutions compared with state-of-the-art. M denotes the number of edges. α is a trade-off parameter between sparsity and fast convergence.

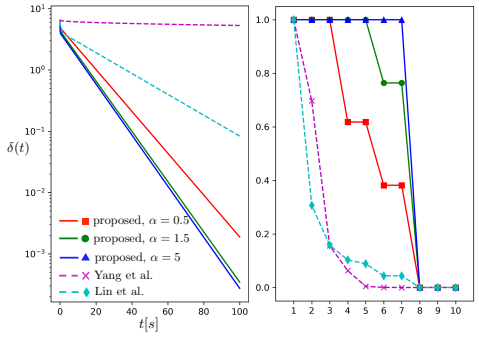

Fig. 2. The convergence of the formation control system (left). $\delta(t)$ is the tracking error exponentially converging to zero. The normalized eigenvalues of the stress matrices (right). The 7th eigenvalue corresponds to the error convergence speed.

References

[1] Z. Li, G. Leus, and R. T. Rajan, "Fast Multiagent Formation Stabilization with Sparse Universally Rigid Frameworks" Submitted to 33rd European Signal Processing Conference (EUSIPCO 2025), in Palermo, Italy.

[2] S. Zhao, "Affine Formation Maneuver Control of Multiagent Systems," in IEEE Transactions on Automatic Control, vol. 63, no. 12, pp. 4140-4155, Dec. 2018, doi: 10.1109/TAC.2018.2798805.

This work is submitted to EUSIPCO 2025 [1]. This work is partially funded by the Sensor AI Lab, under the AI Labs program of Delft University of Technology, and by the European Commission HORIZON-KDT-JU-2023-2-RIA ShapeFuture project.

[3] Z. Lin, L. Wang, Z. Chen, M. Fu and Z. Han, "Necessary and Sufficient Graphical Conditions for Affine Formation Control," in IEEE Transactions on Automatic Control, vol. 61, no. 10, pp. 2877-2891, Oct. 2016, doi: 10.1109/TAC.2015.2504265.

[4] F. Xiao, Q. Yang, X. Zhao and H. Fang, "A Framework for Optimized Topology Design and Leader Selection in Affine Formation Control," in IEEE Robotics and Automation Letters, vol. 7, no. 4, pp. 8627-8634, Oct. 2022, doi: 10.1109/LRA.2022.3188883.

Physics-informed Intelligent Motor Fault Detection on Edge Devices

Sinian Li*, Raj Thilak Rajan*, Edmund Marth†, Patrick Zorn†, Wolfgang Gruber† and Justin Dauwels*
* *Signal Processing Systems, Dept. of Microelectronics, TU Delft, The Netherlands*
† *Institute of Electric Drives and Power Electronics, Johannes Kepler University Linz, Linz, Austria*

Abstract

Intelligent Fault Detection (IFD) has recently gained attention for its ability to enable automatic early detection. Advances in electronic hardware—improved capabilities, lower costs, and miniaturization—now make it feasible to integrate IFD with edge devices for local data analysis and decision-making. While traditional model-based fault detection requires explicit system knowledge, data-driven methods rely solely on training data, demanding large datasets and significant computational power [2]. Given the challenges of complex systems and resource-constrained devices, we propose a low-cost physics-informed (PI) data processing framework for motor commutation angle error detection. Figure 1 shows the overall workflow of the framework, which consists of data collection, data mining, physics-informed feature selection, model training and hardware implementation.

In data mining, notable discrepancies in the averaged torque signal spectra under nominal and faulty conditions reveal characteristic frequency peaks related to electrical and sample-and-hold operation frequencies, which remain consistent across measurements. These peaks guide the proposed feature engineering algorithm, which computes averaged magnitude spectra for both conditions to identify characteristic frequency bins. The resulting set of selected indices (PI features) is then used to extract corresponding rows from the feature matrix for training binary classification models. This PI feature selection identifies the most relevant frequency components, generating a far more effective class separation. Additionally, by reducing feature dimensionality, this approach saves both memory and computational resources. In the inference phase, trained models map the input feature to a discrete class label, predicting the condition of the motor. In this work, we have multiple practical constraints, i.e., highly constrained on-device memory and real-time prediction. Under these conditions, shallow ML models such as Decision Tree (DT), Random Forest (RF), Gradient Boosting (GB), Support Vector Machine (SVM) and a Multiple Layer Perceptron (MLP) models are well-suited. The trained models are then migrated to an Arduino Nano 33 BLE Sense board.

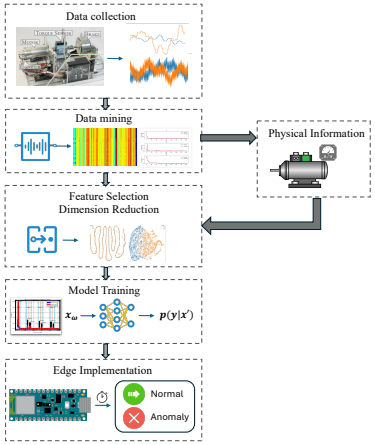

Figure 1. Our proposed physics-informed data processing and model deployment pipeline

The proposed PI features combined with an MLP classification model, achieve the highest test accuracy of 97.8% with the second-lowest memory consumption and inference time, highlighting the potential of tiny ML in IFD even under stringent memory and delay constraints. Notably, the PI features require less than half the dimensionality of conventional statistical features while improving accuracy by 12.7%. Four out of five models exhibit a significant accuracy improvement when incorporating PI features, demonstrating their robustness and generalizability.

In conclusion, we introduced a novel framework, combining physics-based feature extraction with efficient ML models for Edge IFD under stringent memory and delay constraints. Building on these findings, future research will focus on model optimization and compression approaches for more application scenarios.

References

[1] Li, S., et al., "Physics-informed Intelligent Motor Fault Detection for Industrial IoT," , 2025. Submitted to the 33rd European Signal Processing Conference (EUSIPCO 2025), in Palermo, Italy, and unpublished.

[2] Tang, H. D., et al., "A Motor Current Signal-Based Bearing Fault Diagnosis Using Deep Learning and Information Fusion," *IEEE Transactions on Instrumentation and Measurement*, Vol. 69, No. 6, 2020, pp. 3325–3333.

[1] This work is partially submitted to EUSIPCO 2025. It is part of R-PODID project, supported by the Chips Joint Undertaking and its members, including the top-up funding by National Authorities of Italy, Turkey, Portugal, The Netherlands, Czech Republic, Latvia, Greece, and Romania under grant agreement n° 101112338.

Frame-Level Autoregression Yields High-Quality Spatio-Temporal Modelling of Rainfall

Varun Sarathchandran[*1],Ruben Imhof[2], Remko Uijlenhoet[1], Cristian Meo[*1], Justin Dauwels[1]
[1]*Delft University of Technology, The Netherlands.* [2] *Deltares, The Netherlands,* [*] *Equal Contribution.*

Abstract

Forecasting precipitation fields represents a highly complex spatiotemporal modeling task, involving the dynamic interplay of complex spatial structures and rapid temporal evolution in atmospheric data. This study focuses on the Royal Netherlands Meteorological Institute (KNMI) dataset[1], which contains precipitation maps of the Netherlands between 2008 and 2018. Among the latest advancements in the domain, NowcastingGPT [1] achieved state-of-the-art performance on the KNMI dataset, using a video prediction pipeline based on a VQVAE that extracts discrete latents from precipitation maps, and an autoregressive transformer that predicts latent tokens sequentially. However, the autoregression works at a token-level, rather than in time. Such an inductive bias is unideal because the extracted features are highly correlated with each other (e.g., tokens of an image are bidirectionally correlated with each other). Therefore, predicting each token autoregressively is an ill-posed prediction problem, inherently leading to a suboptimal solution. On the other hand, DiffCast [2] introduced a general framework that uses diffusion to model the residual between an internal backbone (e.g., PhyDNet) prediction and the related ground truth, achieving state-of-the-art results on multiple datasets. However, DiffCast requires two-step training, which makes it unclear what kind of dynamics are learned within the model. To address the limitations of prior autoregressive token-based models, we introduce **BlockGPT**- a novel video prediction pipeline which employs an autoregressive transformer that generates one full precipitation map at a time rather than predicting individual tokens sequentially. Our approach, which uses a block attention mask in the transformer, enables bidirectional spatial attention within each frame while maintaining temporal autoregression across frames. This design more naturally aligns with the underlying structure of precipitation data: spatial patterns within a map are best interpreted holistically, while future frames should depend on past context. Furthermore, this architectural shift drastically reduces inference time—our model is $N\times$ faster than token-level autoregressive baselines, where N is the number of latent tokens per map, making it far more practical for real-time nowcasting applications. We evaluate our model against NowcastingGPT and DiffCast. Besides being SOTA, these models represent the two main paradigms in video prediction: discrete, transformer-based generation (NowcastingGPT) and continuous, diffusion-based modelling (DiffCast). Our architecture, which belongs to the former, outperforms both baselines across standard metrics, such as MSE(excluding NowcastingGPT, which we match), MAE, Pearson Correlation Coefficient (PCC), Critical Success Index (CSI), and False Alarm Rate (FAR), demonstrating its effectiveness across both modelling frameworks. Evaluation is performed on a test dataset with 8000 random events. The results, visualised in Figure 1, highlight the strength of our frame-level generation strategy in unifying accuracy, speed, and scalability. Future work will involve evaluating our model on datasets from diverse global regions, encompassing varying climatic conditions and temporal resolutions, to ensure generalizability and assess the model's transferability across different domains.

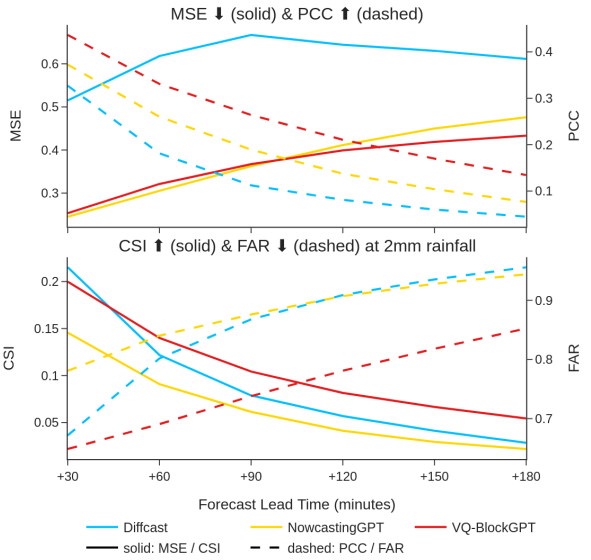

Figure 1. MSE, PCC, CSI, and FAR of BlockGPT and related baselines. BlockGPT achieves the highest PCC, CSI and the lowest FAR, MSE, outperforming all baselines.

References

[1] Meo, C., Roy, A., Lică, M., Yin, J., Che, Z. B., Wang, Y., Imhoff, R., Uijlenhoet, R., and Dauwels, J., "Extreme precipitation nowcasting using transformer-based generative models," *ICLR 2024: Tackling Climate Change with Machine Learning*, 2024.

[2] Yu, D., Li, X., Ye, Y., Zhang, B., Luo, C., Dai, K., Wang, R., and Chen, X., "Diffcast: A unified framework via residual diffusion for precipitation nowcasting," *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

[1]KNMI Dataset:`https://doi.org/10.21944/5c23-p429`

Long-Range Light Messaging: To Backscatter or not to Backscatter?

Tijl Schepens, Gilles Callebaut, and Liesbet Van der Perre
KU Leuven, Ghent, Belgium
tijl.schepens@kuleuven.be

Abstract

The evolution of the Internet of Things (IoT) demands sustainable and energy-efficient solutions capable of long-range communication. Among these, backscatter communication stands out, consuming significantly less power—often 10 to 100 times lower—than active radio solutions. However, traditional limitations of backscatter technology, particularly its restricted communication range, hinder widespread deployment. Recent advancements have introduced Long Range (LoRa)-based backscattering techniques, integrating Chirp Spread Spectrum (CSS) modulation onto low-power, low-complexity tags, thereby extending communication ranges up to hundreds of meters or even kilometers. The question studied in this work is: When to backscatter and when are active transmitters more energy-efficient for long-range and light messaging.

This works systematically reviews existing long-range backscatter technologies and contrasts them with active radio counterparts, focusing on critical aspects such as topology (monostatic vs. bistatic), modulation techniques (On-Off Keying (OOK), Frequency Shift Keying (FSK), CSS), and the use of ambient signals to further minimize power consumption. Practical implementation challenges and solutions—including signal modulation complexity, topology-dependent range improvements, and interference mitigation—are thoroughly assessed. Three included application studies illustrate practical considerations and trade-offs, offering essential insights into choosing between active and backscatter communications for various IoT applications.

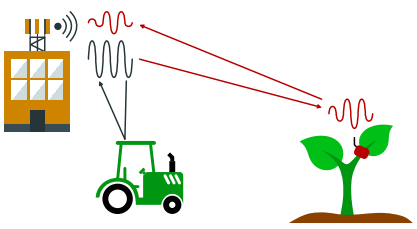

(a) Monostatic backscattering.

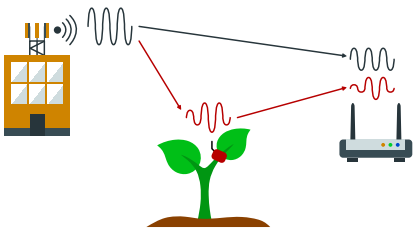

(b) Bistatic backscattering.

Exploring the statistical behavior of Facial Biometrics Based on Neural Embeddings in Privacy-Preserving schemes

Erica Liu
Sig. Proc. Systems, Fac of EE
Eindhoven Univ. of Technology
Eindhoven, the Netherlands
q.liu1@tue.nl

Jean-Paul Linnartz
Sig. Proc. Systems, Fac of EE
Eindhoven Univ. of Technology
Eindhoven, the Netherlands
0000-0002-6098-4043

***Abstract*—In many applications, determining whether a human interacting with a system is the same individual who interacted earlier is a critical task. Facial biometrics have emerged as one of the most prominent modalities of identification in AI-driven systems, yet achieving robust privacy preservation remains a significant challenge. Existing approaches often involve a trade-off, sacrificing either privacy or efficiency. This work addresses the research question whether one can combine AI-driven feature extraction using neural networks with a helper data scheme that preserve privacy by incorporating a one-way hash function and maintain identification accuracy. A concatenation may be straightforward if feature vectors act as independent Gaussian random variables. To verify this, we evaluate an identification capacity metric, grounded in information theory, as a step towards a quantitative evaluation of the (potential) performance of AI-driven identification systems.**

***Index Terms*—identification system, privacy-preserving, capacity**

I. Introduction

In modern intelligent systems, detecting and identifying individuals is essential for building smart cities. Person re-identification (Re-ID) plays a pivotal role in scenarios where individuals need to be tracked across different cameras or scenes [1]. For instance, consider the case of cyclists navigating urban environments: Re-ID systems can monitor their movements across multiple intersections, ensuring safety, optimizing traffic flow, and supporting data-driven urban planning [2]. Similarly, face recognition systems are ideal for identifying individuals in scenarios where only facial images are captured, such as at entry checkpoints [3]. Although they use various biometrics, the underlying requirements for robust feature extraction and privacy-preservation share significant overlap [4].

This work introduces a novel evaluation metric for privacy-preserving feature extraction, validated on face recognition tasks due to their controlled environment and established benchmarks. Although initially applied in face recognition, specifically ArcFace in this work [5], the principles underpinning this metric—balancing accuracy, privacy, and efficiency—are directly extendable to Re-ID scenarios like cyclist monitoring. Future iterations will refine this approach for body-oriented features, bridging the gap between these complementary tasks.

As shown extensively in literature, e.g. [6], biometrics can be used to achieve identification and face recognition of individuals. Conventional methods tend to perform statistical analysis on collected biometric data. Specifically, traditional biometric identification systems aim to minimize intra-class variance through robust and invariant feature extraction, while simultaneously maximizing inter-class variance to ensure clear separation between different identities in the feature space [7].

This paper studies a particular class of Re-ID in which privacy is important. That is, the system needs to be capable of detecting whether a person observed by a camera has been seen before, without identifying that person specifically. In the previous example of traffic monitoring, it is important that the system can attribute camera observations to traffic participants and follow their flow as living probes, but there is no need, or even a prohibition, to know or to recognize the identity of these probe participants. Preferably, we preserve and guarantee privacy by designing the system such that recognition of specific persons is intrinsically hard, even for an inside attacker who has access to data stored by the system. In our example: ”It took a reference cyclist who adheres to the traffic regulations 15 minutes to travel from the railway station to the church.” But not: ”the pastor ignores a red traffic light occasionally” (recognizing the person) nor ”The person who entered the church at 10:45 is a black female” (gaining side information).

When privacy is a concern, biometric features are often further transformed using privacy-preserving techniques, such as coarse representations. Our work is motivated by the observation that privacy-preserving techniques are known in information theory, allowing cryptographic transformations such as hash functions to secure sensitive information while maintaining identification performance [8]. However, the combination of these with AI techniques seems to have received limited attention. Recently, with the development of deep learning in image processing, the information in biometrics can be leveraged more effectively. For instance, deep learning models

can learn representative features from images, capturing both detailed and global characteristics of an image (person) [9].

There are still several problems in this field. The first is accurately (re-) identifying a person anonymously, thus maintaining the individual privacy. The second one is designing and deploying the identification system on the resource-constrained devices to achieve real-time efficiency [10]. In addition, considering the privacy requirement, accurate feature extraction models in various environments should be a precondition. Based on the more accurate model, we not only exploit the capacity better in seeking for good trade-off between accuracy and efficiency, but also to make a better trade-off between accuracy and privacy. Helper data schemes (HDS) were originally designed to work with random variables, in theoretical models often jointly Gaussian, independent and with the same variance, to protect privacy. Capacity as a classic information theory term refers to the maximum amount of information that can be sent to any channel or medium. Willems et al. [11] introduced the concept of capacity into the identification system, to evaluate the number of individuals that the system can be identified error-free. We use these ideas to explore the suitability of AI generated feature vector to apply helper data schemes

In order to investigate to what extent the progress in AI pre-processing (NN-based) for face recognition can effectively be used in combination with helper data schemes (HDS), we study a number of aspects. The subsequent sections review the requirements for privacy in HDS and analyze the feature embeddings extracted from face datasets of various individuals. Based on feature embeddings, the inter-class (between) and intra-class (within) variances can be estimated from feature vectors to prepare the required information for HDS [12]. Additionally, we estimate the capacity to check if our combination leads to theoretically reasonable results. Rather than presenting specific numerical capacity estimates in the introduction, we note that our analysis indicates that the theoretical identification capacity based on Gaussian assumptions is lower than expected from the networks' performance. Detailed quantitative evaluations and discussions are provided in later sections.

II. Related work

Biometric identification, such as face recognition and person re-identification (Re-ID), is important in many smart systems. A key challenge is to maintain high identification accuracy while protecting user privacy. Traditional methods often involve trade-offs between privacy and performance. Recent works combine deep feature extraction with privacy-preserving techniques [13]. Other studies propose information-theoretic tools to evaluate identification systems [14].

A. Deep Learning in Person Re-Identification

Deep learning significantly enhances person Re-ID by learning features that reduce intra-class variance and increase inter-class separation. Ye et al. [15] and Ming et al. [16] provide comprehensive surveys on deep learning for Re-ID, addressing scalability and deployment in constrained environments. However, existing deep learning-based Re-ID methods often overlook the potential risks of privacy leakage inherent in their design and deployment, particularly with respect to the storage of feature embeddings in such systems.

ArcFace, introduced by Deng et al. [17], is a widely used deep face recognition model employing an additive angular margin loss. Although originally trained for face recognition, ArcFace's embeddings generalize well to other identity recognition tasks. In this work, we adopt ArcFace as a pre-trained feature extractor to generate identity embeddings, which are then processed using privacy-preserving mechanisms.

B. Privacy-Preserving Techniques in Biometric Systems

Various methods have been developed to enhance privacy in biometric systems while maintaining recognition performance.

Differential Privacy: Chamikara et al. [18] proposed PEEP, which adds noise to facial features before storage or matching. This preserves privacy and maintains an accuracy range of 70% to 90%. However, from an information-theoretic perspective, these noise templates still retain significant mutual information with the original data, and any noise addition or blurring can negatively impact re-identification performance.

Homomorphic Encryption: Face templates can be encrypted such that comparisons are performed directly in the encrypted domain. Additional optimizations like clustering have been introduced to reduce computational costs [19]. However, while homomorphic encryption enables secure computation, it does not eliminate the need for decryption; thus, the overall system security still critically depends on safeguarding the decryption key—leaving it potentially vulnerable to insider threats [20].

De-Identification: Dou et al. [21] presented Person Identity Shift (PIS), which alters visible identities in images while preserving relative identity relationships. This enables privacy-preserving data usage during training.

Helper Data Scheme with Hashing: Helper Data Schemes (HDS) are cryptographic tools for extracting reliable secrets from noisy biometric inputs. During enrollment, helper data and a hash of the original biometric are stored. At authentication, the helper data enables reconstruction of the biometric feature for hash verification. This process aligns with information-theoretic views of communication over noisy channels [22], [23]. HDS prevents direct recovery of the biometric template, making it suitable for privacy-preserving Re-ID systems. However, we are not aware of any studies that have explored the use of Helper Data Schemes together with state-of-the-art neural networks for face recognition.

C. Information-Theoretic Evaluation Metrics

Information theory provides a framework to assess the reliability of identification systems. Identification capacity measures the maximum number of individuals a system can distinguish without error, which is crucial for resource-limited environments such as IoT edge devices. Willems et al. [11]

introduced this concept for biometric systems, showing how capacity can serve as a theoretical limit for secure and accurate identification.

III. System description and formulation of an evaluation model

A. Data preprocessing

Face images from various persons consists of our dataset. Additionally, the face images should be aligned before the feature extraction so we only extract features from those faces instead of other parts of persons.

We apply the VGGFace2 and finish the face alignment using Scipy, and we controlled the number of images per person. Finally, we got 42 persons in the dataset, each person has 300 face images(RGB format reserved).

B. Feature extraction

We extract features from those pre-processed face images applying the deep convolutional neural networks. Specifically, we firstly use the encoder of Variantional AutoEncoder (VAE) due to its capacity to notice the global and local parts for images. Secondly, we select the pre-trained ArcFace as our feature extractor since it achieves highest performance for several classic benchmarks in face recognition. The ArcFace can capture the facial characteristics well and the researcher use the Arc loss to enhance the discriminative power for various person faces. In this case, the model can distinguish persons better than using previous softmax function as the loss function. The ArcFace is pre-trained on MS1MV2 dataset, which can extract 512 features for one image, following with a similarity calculation we can re-identify individuals. We apply this feature extraction on the entire dataset to obtain the features for analysis.

C. Enrollment (or first identification)

At enrollment, we have a total of N dimensions and we index the dimension $n = 1, 2, \ldots, N$. We have M individuals, whom we index $m = 1, 2, \ldots, M$, for each individual, we have S samples in the dataset, and we index them $s = 1, 2, \ldots, S$. The person's face images $\vec{X}(m)$ contains variability(light variability, angel variability, etc.), specifically enrollment face image, is processed by the feature extractor and principal component analysis (PCA), yielding feature vector $\vec{x}_{m,s}$ for each sample. The reason we applied PCA is to produce more independent features after projection. We also analyze the statistical properties of extracted features without PCA to calculate how can we determine the principal components with a selected explained variance ratio.

In the analysis of the variability among dimensions and individuals, we fix individual m and compare the distributions of $X_1(m), \ldots, X_N(m)$ from different dimensions; and for the fixed dimension n, we compare the distributions of $X_n(1), \ldots, X_n(M)$ from various individuals.

Furthermore, we analyze the intra/inter-class variation in enrollment. Since we have S observations for M enrolled individuals, the mean values can be estimated from taking averages of samples:

$$\begin{aligned} \vec{\mu}_m &= \frac{1}{S}\sum_{s=1}^{S}\vec{x}_{m,s} \\ \vec{\mu} &= \frac{1}{M}\sum_{m=1}^{M}\vec{\mu}_m \end{aligned} \tag{1}$$

We estimate the intra-/inter-class within/between) variance in the usual way

$$\begin{aligned} \Sigma_w &= \frac{1}{M}\sum_{m=1}^{M}\left[\frac{1}{S-1}\sum_{s=1}^{S}\left(\vec{x}_{m,s}-\vec{\mu}_m\right)\left(\vec{x}_{m,s}-\vec{\mu}_m\right)^T\right] \\ \Sigma_b &= \frac{1}{M-1}\sum_{m=1}^{M}\left(\vec{\mu}_m-\vec{\mu}\right)\left(\vec{\mu}_m-\vec{\mu}\right)^T \end{aligned} \tag{2}$$

in which $\vec{x}_{m,s}$, $\vec{\mu}_m$, $\vec{\mu}$ are the vectors over dimension N. The accuracy of these estimates depends on the number of samples M, and on the extent of correlation among them. An imbalanced dataset may lead to underestimation of both inter-class and intra-class variances, which in turn can result in a lower estimated identification capacity.

D. (Re-) Identification

In the identification phase, we combined the foreknowledge in ArcFace. ArcFace calculated the cosine similarity between the feature vectors and set a threshold to determine if two images are from the same person. The key point is that similar images can produce similar directions in latent space, which is the feature vectors after the extraction from images. Instead of keeping the original feature vectors, we convert the feature vectors to unit vectors and compare the values of each dimension to the original registered values of each dimension. If the difference in all dimensions is less than the threshold value for that dimension, then we can determine which person the test image belongs to. This is also why we need to analyze the distributions of variables to check how to set our thresholds for dimensions.

Similarly with in enrolment phase, person's face image $\vec{X}(m)$ containing noise in identification phase is also processed by the feature extractor and projected using the same projection matrix as in enrolment, yielding feature vector $\vec{x}_{m,s}$ with dimension N.

E. Capacity evaluation

The capacity in biometric identification is a measurement of the amount of individuals that can be identified without errors [11]. Hence, we estimate the theoretical capacity of our biometrics and compare it to the actual capacity in the real-world scenarios.

Based on the biometric model, there are various hidden biometric random variable $R \in \mathbb{R}^N$ for each individual and the features we obtain from enrolment and identification contains independent additive Gaussian variations N_e and N_i as shown

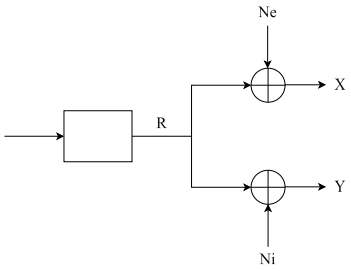

Fig. 1. Biometric model for each individual

in fig 1, the samples we observed in enrolment and identification phase are X and Y. As explored in previous research [24], the model in biometric identification systems, referred to as the 'within-/between-class distribution' with variances σ_w^2 and σ_b^2, accounts for variations within individuals and between individuals to achieve better identification performance.

We anticipate the variation of one individual can be reduced averaging all samples of an individual to produce the optimal estimation of hidden biometric value $R \approx \vec{\mu}$. And the theoretical capacity for our biometrics can be estimated by the mutual information between the enrolment and identification samples as

$$I(X;Y) = h(X) + h(Y) - h(X,Y) \tag{3}$$

in which $X \sim \mathcal{N}(\mu, \Sigma_X)$ and $h(X,Y)$ is the joint entropy, $Y \sim \mathcal{N}(\mu, \Sigma_Y)$, also, since N_e and N_i are independent additive Gaussian, we have $\Sigma_X = \Sigma_R + \Sigma_{N_e}$ and $\Sigma_Y = \Sigma_R + \Sigma_{N_i}$. The h represents the entropy function, and due to the Gaussian variables assumption, we can get

$$\begin{aligned} I(X;Y) =& \frac{1}{2}\log_2\left((2\pi e)^N |\Sigma_X|\right) + \frac{1}{2}\log_2\left((2\pi e)^N |\Sigma_Y|\right) \\ & - \frac{1}{2}\log_2\left((2\pi e)^{2N} |\Sigma_{XY}|\right) \\ =& \frac{1}{2}\log_2 \frac{|\Sigma_X|\,|\Sigma_Y|}{|\Sigma_{XY}|} \end{aligned} \tag{4}$$

In the above equation, Σ_{XY} can be calculated as equation 5.

$$\begin{aligned} \Sigma_{XY} &= \begin{bmatrix} \mathbb{E}\left[X^2\right] - \mathbb{E}[X]^2 & \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y] \\ \mathbb{E}[YX] - \mathbb{E}[Y]\mathbb{E}[X] & \mathbb{E}\left[Y^2\right] - \mathbb{E}[Y]^2 \end{bmatrix} \\ &= \begin{bmatrix} \Sigma_R + \Sigma_{N_e} & \Sigma_R \\ \Sigma_R & \Sigma_R + \Sigma_{N_i} \end{bmatrix} \end{aligned} \tag{5}$$

Or it can be calculated according to the property of joint covariance matrix as follows,

$$\begin{aligned} \Sigma_{XY} &= \begin{bmatrix} K_{XX} & K_{XY} \\ K_{YX} & K_{YY} \end{bmatrix} \\ &= \begin{bmatrix} \Sigma_X & \Sigma_{XY} \\ \Sigma_{XY} & \Sigma_Y \end{bmatrix} \\ &= \begin{bmatrix} \Sigma_R + \Sigma_{N_e} & \Sigma_R \\ \Sigma_R & \Sigma_R + \Sigma_{N_i} \end{bmatrix} \end{aligned} \tag{6}$$

The Σ_{XY} is the joint covariance matrix because of the joint entropy $h(X,Y)$, and N is the vector length(dimensions in our case), the $|\Sigma|$ is the determinant of covariance matrix. hence Σ_{XY} is a matrix in $\mathbb{R}^{2N\times 2N}$, while Σ_X and Σ_Y are matrices in $\mathbb{R}^{N\times N}$.

We can calculate the $|\Sigma_X|$ and $|\Sigma_Y|$ using the product of eigenvalues, however, the $|\Sigma_{XY}|$ has various size with previous two terms, leading to the difficulty of understanding the relationship between capacity and the number of biometric features. Thus, we convert it to equation 7 based on the invertible property.

$$\begin{aligned} |\Sigma_{XY}| = \big|&(\Sigma_R + \Sigma_{N_e})(\Sigma_R + \Sigma_{N_i}) \\ &-\Sigma_R (\Sigma_R + \Sigma_{N_i})^{-1} \Sigma_R (\Sigma_R + \Sigma_{N_i})\big| \end{aligned} \tag{7}$$

With this conversion, we can transform equation 4 to the following calculation

$$\begin{aligned} I(X;Y) &= \frac{1}{2}\log_2 \frac{\prod_{i=1}^{N} (\lambda_X)_i (\lambda_Y)_i}{\prod_{i=1}^{N} (\lambda_{XY})_i} \\ &= \frac{1}{2}\sum_{i=1}^{N} \log_2 \frac{(\lambda_X)_i (\lambda_Y)_i}{(\lambda_{XY})_i} \end{aligned} \tag{8}$$

in which the λ represents the eigenvalue and $i \in \{1..N\}$ indexes the dimension. Regrettably, this eigenvalue-based calculation does not yield plausible and reliable results. One reason is that the number of samples affects the estimation of Σ_w and Σ_b, which in turn influences the corresponding eigenvalues. Another reason is that many dimensions among neural network-extracted features are highly correlated and add only a little additional information. As a result, the values of λ_X, λ_Y, λ_{XY} are sometimes closed to zero, which can lead to nonsensical outcomes when applying the logarithm function.

Based on the estimation of $R \approx \vec{\mu}$, we calculate the covariance and obtain an approximation of the biometric covariance, i.e.,$\Sigma_R = \Sigma_b$. Similarly, the noise covariances can be represented as $\Sigma_{N_e} \approx \Sigma_{N_i} \approx \Sigma_w$. Therefore, we can calculate the theoretical capacity of the biometrics based on the intra-/inter-class covariance matrices as in 9. However, a similar issue arises here: the covariance matrices can be ill-conditioned, leading to computational difficulties when applying the logarithm function.

$$I(X;Y) = \frac{1}{2}\log_2\left(1 + \frac{|\Sigma_b|^2}{|\Sigma_w|^2 + 2\,|\Sigma_w|\,|\Sigma_b|}\right) \tag{9}$$

To avoid these numerical problems, we ignore the correlation between feature dimensions, that is we simplify the covariance matrix $|\Sigma_b|$, $|\Sigma_w|$ as a diagonal matrix. Then we can directly apply the diagonal values to calculate the capacity according to equation 9. Then the capacity expression is shown as in equation 10, in which $(\sigma_b)_i, (\sigma_w)_i$ represent the diagonal value of $|\Sigma_b|$, $|\Sigma_w|$ in $i-th$ dimension. In this simplified scenario, the capacity estimates are based on

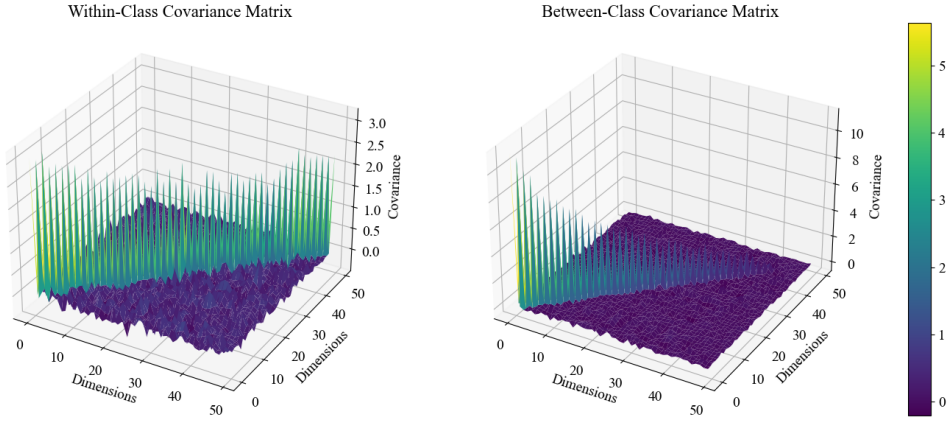

Fig. 2. The within/between-class covariance matrix (only show the first 50 dimensions)

redundant information, resulting in values that underestimate the actual capacity.

$$I(X;Y) = \frac{1}{2}\sum_{i=1}^{N}\log_2\left(1+\frac{(\sigma_b)_i^2}{(\sigma_w)_i^2+2\,(\sigma_w)_i\,(\sigma_b)_i}\right) \quad (10)$$

Note that equality holds if the dimensions are independent.

IV. Experimental Analysis

As mentioned previously, the overview process of privacy-preserving re-ID is that, we firstly extract representative features from images in the dataset; We can apply quantization to produce the coarse representation for features, and the helper data scheme can lead to the unique and reproducible template for each person.

However, since the properties of feature vectors are important in the privacy-preserving re-ID, we firstly test the receiver operating characteristic (ROC) curve to make sure the pre-trained ArcFace performs well in non-privacy preserving face recognition task. Besides, we get the features statistical analysis to check if they are the expected independent Gaussian distribution, and to check the distribution difference among people and dimensions.

Furthermore, if the feature statistical analysis is not the independent as expected, we try principle component analysis to force the features to be independent. Then we get the capacity curve according to the equations in section 3. After all these features statistical analysis, we apply the helper data scheme to achieve the privacy-preserving for the biometric identification system, in which quantization followed by one-way hash function is required to have proper parameters, such as intervals.

V. Results

A. Features statistical analysis

As stated in the system description, we analyze the distribution of each feature dimension after the feature extraction. We first check the cumulative distribution functions (CDFs) and probability density functions (PDFs) to analyze the statistical property of features $x_{i=1}^{N}$. When we fixed the person identity, we observe the differences in both the CDFs and the PDFs across different feature dimensions as shown in fig 3 and fig 4. While the CDFs and PDFs generally resemble a Gaussian shape, variations in location and steepness across dimensions indicate differences in mean and variance. When we fixed the dimension, we compare the distributions of various individuals in the same dimension to check the distinguish-ability. Part of the results shows in fig 5 and fig 6. The results show part of the distinguish ability of features among various individuals.

Moreover, to assess the overall redundancy in our 512-

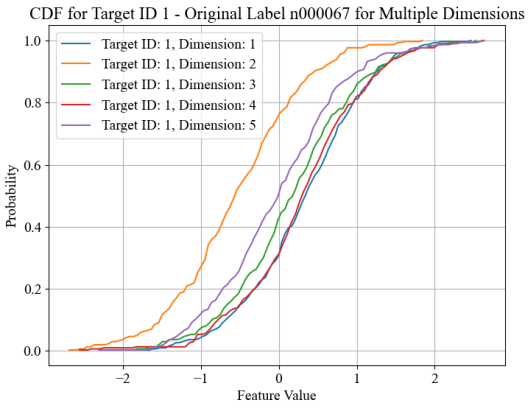

Fig. 3. The CDFs of $X_1(1), \ldots, X_5(1)$ without PCA

dimensional embedding space, we have the population-wide correlation heatmap. The results shows that while the majority of feature dimensions are largely uncorrelated, we identify

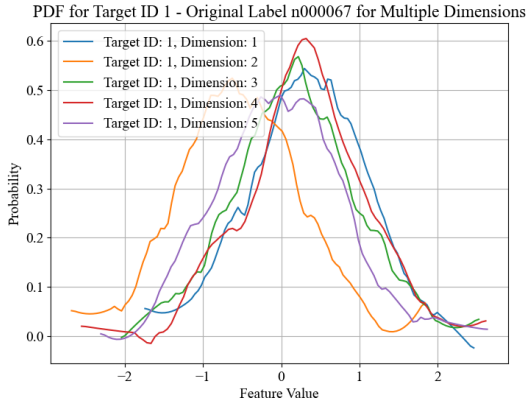

Fig. 4. The PDFs of $X_1(1), \ldots, X_5(1)$ without PCA

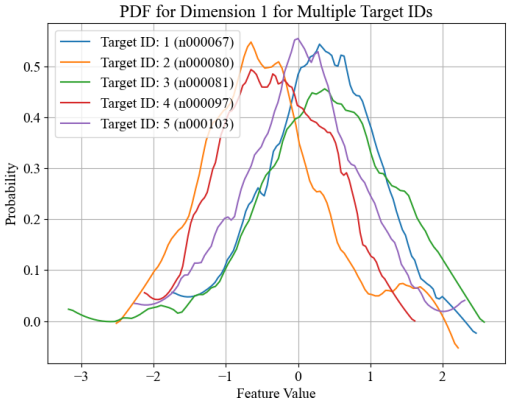

Fig. 6. The PDFs of $X_1(1), \ldots, X_1(5)$ without PCA

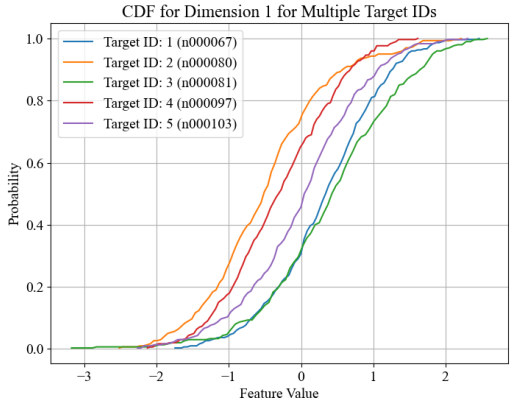

Fig. 5. The CDFs of $X_1(1), \ldots, X_1(5)$ without PCA

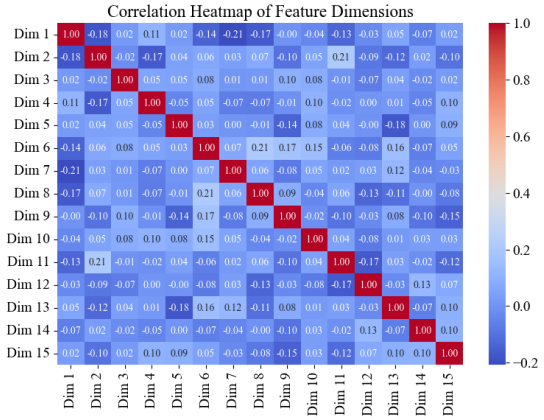

Fig. 7. The population-wide correlation heatmap for the first 15 dimensions

dozens of dimension pairs ($|r| > 0.3$) exhibiting moderate intercorrelation, suggesting partial redundancy in the learned representation. The first 15 dimensions correlation is shown in fig 7.

Additionally, we tried principle component analysis (PCA) to check the effective dimensionality and variance distribution of the feature embeddings. Figure 8 presents the full eigenvalue spectrum (left) and per-component explained-variance ratio (right) for the 512-dimensional enrollment covariance matrix of feature embeddings. The left panel shows that the largest eigenvalue reaches approximately 14 but decays sharply to below 1 by the 50th component and approaches zero around the 300th. The right panel plots the individual explained-variance ratio (blue) and cumulative variance (orange), revealing that the first 10 components capture only about 60% of the total variance, the first 50 about 70%, and roughly 200 components are required to explain 95%. This pattern indicates that, while the leading principal axes concentrate much of the between-class signal, a substantial tail of mid-ranked dimensions still carries non-negligible identity information, thereby limiting the extent of safe dimensionality reduction.

B. Inter/Intra-class variance

Additionally, we analyze the within group variance, which can capture the variability of different images from one person. Also, the between group variances are calculated according to section 3. And the results are shown in figure 2. We can see that the within-class variances (left) are relatively modest and fairly uniform after Dim 10, reflecting consistent, low-level fluctuations of each feature within an individual's samples. In contrast, the between-class variances (right) exhibit larger values in the earliest dimensions and then decay rapidly toward zero. This pattern indicates that the first few embedding channels possess greater discriminative power.

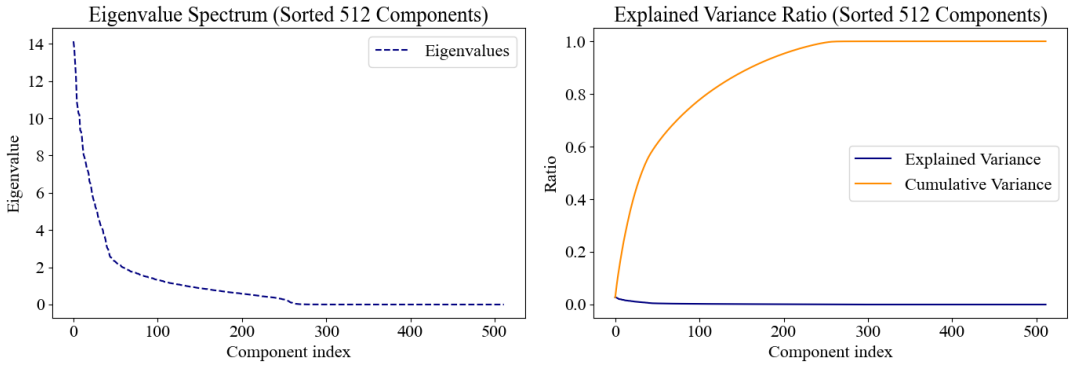

Fig. 8. Eigenvalue spectrum and explained-variance ratio for the 512-dimensional enrollment covariance matrix

Furthermore, figure 9 visualize the inter/intra-class variance for two embedding dimensions. Dashed curves represent the population-wide distributions (inter-class variance) for dimension 1 and dimension 2, while solid curves depict the single-subject distributions (intra-class variance) for the same dimensions. The differing overlaps and peak locations indicate that the degree of separation between inter- and intra-class variance varies by dimension, with dimension 1 exhibiting greater class-level dispersion relative to individual-level variation than dimension 2.

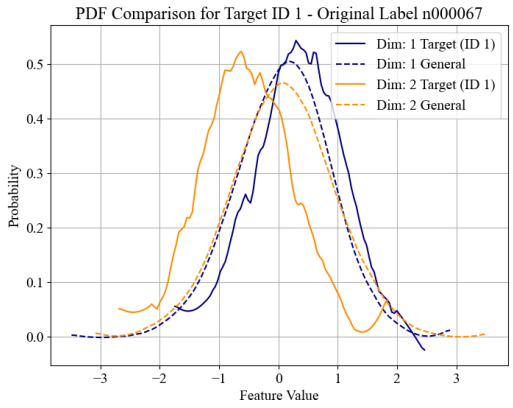

Fig. 9. The PDFs of population-wide ID and single ID (Inter- and intra-class variance)

C. ROC Curve

To verify the identification power of our embeddings before applying any privacy-preserving techniques, we computed the ROC curve using the raw feature vectors, which are the direct outputs of a pre-trained ArcFace model without further processing. Figure 10 shows the ROC curve for the raw ArcFace feature vectors, yielding an AUC of 0.93 and demonstrating strong separability between genuine and impostor pairs. These points highlight the trade-off between security (low FPR) and usability (high TPR), indicating that the decision threshold can be tuned to satisfy different identification requirements.

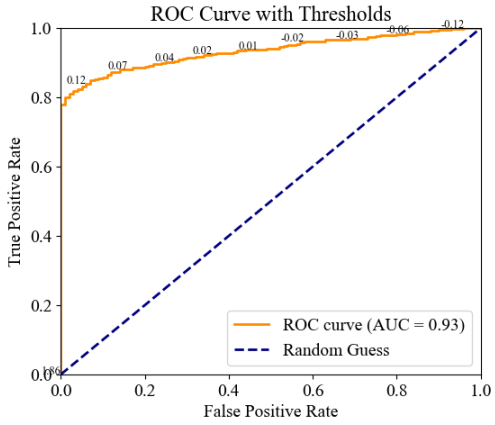

Fig. 10. The ROC curve with thresholds

D. Capacity Estimation

As stated in our previous section, we can estimate the capacity across feature dimensions based on equation 10. Figure 11 shows the identification capacity for the first 50 embedding dimensions, plotted in their original channel order. Capacity falls sharply from approximately 0.7 in dimension 1 to near zero by dimension 40; dimensions beyond 50 are omitted because their capacity is effectively zero. This trend indicates that the earliest dimensions capture the bulk of between-class variance, while later dimensions contribute progressively less discriminative power.

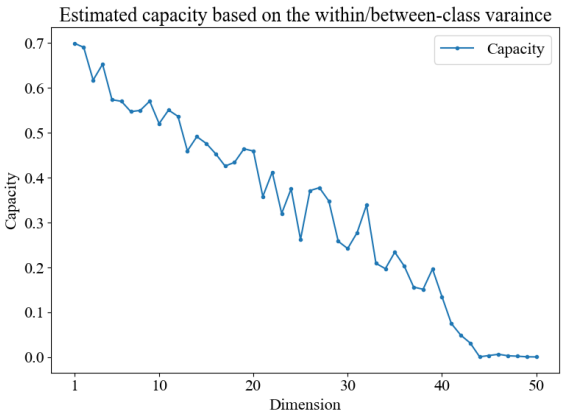

Fig. 11. The estimated capacity over the first 50 dimensions (Ignoring correlations between dimensions)

VI. Conclusion

In summary, our work explored the feasibility of combining AI-driven feature extraction with a helper data scheme that can allow a one-way hash function. The intention is to maintain identification accuracy while preserving privacy in a cryptographic way. Our analysis reveals that the feature vectors over dimensions seem to approximate a Gaussian behavior, as both the PDF and CDF closely resemble a Gaussian distribution, and dimensions are correlated to some extent. This correlation, along with the observed variance patterns, plays a critical role in identification performance. Nonetheless, we observed a number of challenges. If we calculate capacity from a theoretical Gaussian model, even if we optimistically assume independent features is lower than expected. We extracted statistical parameters by estimating covariances as root mean square errors in a limited data set. It appeared that using these in capacity formulas results in numerical problems and unrealistic values. Specifically, the within-class variance did not exhibit the minimal differences anticipated for a single individual, while the between-group variance analysis indicates that lower-dimensional representations reveal a more pronounced inter-individual variability. Moreover, the identification capacity metric shows a decline as dimensionality increases, suggesting that lower-dimensional representations preserve more relevant information for enhancing identification ability. In the further development of the system, this will inevitably present challenges in properly setting quantization intervals to prepare the feature values for cryptographic operations, such as one-way hashing. Overall, our findings contribute toward establishing the viability of combining neural network-based feature extraction with helper data schemes. This approach offers a path forward for achieving privacy preservation without significantly compromising the discriminative power of AI-based biometric systems.

References

[1] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. Hoi, "Deep learning for person re-identification: A survey and outlook," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 6, pp. 2872–2893, 2021.

[2] X. Liu, R. Layne, T. M. Hospedales, and S. Gong, "A deep learning-based approach to progressive vehicle re-identification for urban surveillance," in *European Conference on Computer Vision*, pp. 869–884, Springer, 2016.

[3] A. Unknown, "A comprehensive overview of biometric recognition in smart cities," *Unknown Source*, Recent. Retrieved from Semantic Scholar.

[4] E. Wenger, A. N. Bhagoji, M. Chiang, and P. M. Li, "Feature vector stealing attacks against deep facial recognition models," *arXiv preprint arXiv:2202.05760*, 2022.

[5] Q. Dang, Q.-V. Ha, D. Nguyen, and C. Yuen, "Smart noise: A privacy-preserving approach for learning deep face representations," in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3153–3157, IEEE, 2022.

[6] N. V. Boulgouris, K. N. Plataniotis, and E. Micheli-Tzanakou, *Biometrics: theory, methods, and applications*. John Wiley & Sons, 2009.

[7] S. Gosavi and P. Mohod, "Evaluation of inter-class and intra-class variance for face recognition using deep convolutional neural network," in *2017 International Conference on Intelligent Computing and Control (I2C2)*, pp. 1–4, IEEE, 2017.

[8] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," in *International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 523–540, Springer, 2008.

[9] H. Lin, K. Zhang, Y. Wu, R. Zhang, Y. Ji, and S. Yan, "Gaitset: Regarding gait as a set for cross-view gait recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 11669–11676, 2020.

[10] A. Zahra, N. Perwaiz, M. Shahzad, and M. M. Fraz, "Person re-identification: A retrospective on domain specific open challenges and future trends," *Pattern Recognition*, p. 109669, 2023.

[11] F. Willems, T. Kalker, J. Goseling, and J.-P. Linnartz, "On the capacity of a biometrical identification system," in *IEEE International Symposium on Information Theory*, pp. 82–82, 2003.

[12] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, M. Li, Y. Lin, and Y. Mu, "Softmax loss and its variants: A survey," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, IEEE, 2019.

[13] C. Rathgeb, K. Raja, R. Raghavendra, and C. Busch, "Deep face recognition: A survey," *ACM Computing Surveys (CSUR)*, vol. 54, no. 5, pp. 1–35, 2021.

[14] C. Rathgeb and C. Busch, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP Journal on Information Security*, vol. 2011, no. 1, pp. 1–25, 2011.

[15] M. Ye, J. Shen, G. Lin, T. Xiang, and L. Shao, "Deep learning for person re-identification: A survey and outlook," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2022.

[16] Z. Ming *et al.*, "Deep learning for person re-identification: A survey and outlook," *arXiv preprint arXiv:2110.04764*, 2021.

[17] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4690–4699, 2019.

[18] M. Chamikara, E. Bertino, D. Liu, and S. A. Camtepe, "Peep: Privacy-enhanced yet efficient and accurate biometric identification using differential privacy," *arXiv preprint arXiv:2005.10486*, 2020.

[19] S. Song, J. Lee, and T. Kim, "Privacy-preserving face recognition using homomorphic encryption and clustering-based optimization," *IEEE Transactions on Information Forensics and Security*, 2025. To appear.

[20] J. H. Cheon, M. Kim, Y. Lee, and Y. Song, "A practical multiparty computation protocol for privacy-preserving face recognition," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp. 553–567, ACM, 2019.

[21] Z. Dou, Z. Zheng, L. Zhang, and T. Xiang, "Towards privacy-preserving person re-identification via person identity shift," *arXiv preprint arXiv:2207.07311*, 2022.

[22] T. Ignatenko and F. M. Willems, "Biometric security from an information-theoretical perspective," *Foundations and Trends in Com-*

munications and Information Theory, vol. 11, no. 3-4, pp. 209–420, 2015.

[23] J. Merkle and B. Tams, "Quantization in zero leakage helper data schemes," *EURASIP Journal on Advances in Signal Processing*, vol. 2016, no. 1, pp. 1–13, 2016.

[24] E. J. Kelkboom, G. G. Molina, J. Breebaart, R. N. Veldhuis, T. A. Kevenaar, and W. Jonker, "Binary biometrics: An analytic framework to estimate the performance curves under gaussian assumption," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 40, no. 3, pp. 555–571, 2010.

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

Graph Topology Identification Based on Covariance Matching

Yongsheng Han, Alberto Natali, Geert Leus
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Delft, The Netherlands
{y.han, a.natali, g.j.t.leus}@tudelft.nl

***Abstract*—This paper addresses graph topology identification for applications where the underlying structure of systems like brain and social networks is not directly observable. Traditional approaches based on signal matching and spectral templates have limitations, particularly in handling scale issues and sparsity assumptions. We introduce a novel covariance matching methodology that efficiently reconstructs the graph topology using observable data. For the structural equation model (SEM) using an undirected graph, we demonstrate that our method can converge to the correct result under relatively soft conditions. Furthermore, we extend our methodology to polynomial models and any known distribution of latent variables, broadening its applicability and utility in diverse graph-based systems.**

***Index Terms*—graph topology identification, covariance matching, structural equation model, polynomial model**

I. Introduction

Graph topology identification remains a critical issue in graph signal processing (GSP), where systems are modeled as networks, yet their actual underlying structure is often invisible. Examples of such systems include brain functional connectivity networks and social networks. In these applications, while the direct graph structure is not observable, nodal data is typically available. For instance, in academic networks [1], the advisor-advisee links may not be visible, yet we can analyze collaborative patterns to uncover these connections. Similarly, in brain networks [2], neural signals provide indirect clues about the connectivity. Therefore, the primary challenge in graph topology identification lies in deducing the hidden graph structure from these nodal observations, a task that is fundamental for analyzing and understanding the interactions within these networks.

The structural equation model (SEM) is a popular tool to link nodal data with the graph, and it has been frequently used in graph topology identification [3], [4], [5], [6]. A remarkable result was provided by [6], where it was demonstrated that for sparse directed acyclic graphs (DAGs), the graph can be uniquely determined when the unknown external inputs are Gaussian with equal variance. This foundational work spurred further developments, leading to more efficient algorithms as evidenced by [7] and [8]. However, for undirected graphs, methods using signal matching have performed poorly without the presence of exogenous variables, even with the introduction of sparsity constraints [9], [10]. On the other hand, spectral template-based approaches, such as the polynomial graphical lasso (PGL) algorithm, have shown potential in handling certain graphs [11], but the results typically differ from the true structure by a scale factor and require extensive restrictions (sparsity and sign) on the graph for the method to be effective.

To the best of our knowledge, no existing work has exhaustively addressed topology identification for undirected graphs using a SEM. This paper fills this gap by proposing a novel covariance matching-based method. We will prove that under relatively soft conditions, our proposed method consistently converges to the correct result without encountering the scale issue often associated with other approaches, and our method is more robust. Furthermore, we do not need to make any assumptions about sparsity. We will also extend our method to accommodate polynomial models and any known distribution of latent variables, broadening the applicability of our approach in graph topology identification.

*This work is partially supported by the NWO OTP GraSPA proposal #19497, and the EU HORIZON-CHIPS-JU-2023-2-RIA ShapeFuture project, under grant agreement No 101139996.

II. Preliminaries

In this section, some background information is presented that is required to explain the main contributions of this work.

Graph signal processing: We describe a graph as $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{S}\}$, where $\mathcal{V} = \{1, \dots, N\}$ represents the set of vertices, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ denotes the set of edges, and $\mathbf{S} \in \mathbb{R}^{N \times N}$ is referred to as the graph shift operator (GSO). It is an $N \times N$ matrix that captures the local structure of the graph, whose entries S_{ij} are non-zero only if $i = j$ or if $(j, i) \in \mathcal{E}$ [9]. The adjacency matrix or the combinatorial graph Laplacian are matrices that can be used as GSO. Each node $i \in \mathcal{V}$ is associated with a scalar value x_i. By stacking these values into a vector $\mathbf{x} = [x_1, \dots, x_N]^\top \in \mathbb{R}^N$, we obtain what is known as a graph signal.

Structural equation model: Considering a simplified SEM excluding the influence of exogenous variables, the internal structure of a graph signal $\mathbf{x}$ can be expressed as

$$\mathbf{x} = \mathbf{S}\mathbf{x} + \mathbf{e}. \tag{1}$$

Here, the matrix $\mathbf{S}$ represents the GSO of the graph. Clearly, this GSO should have zero diagonal entries indicating no self-influence and hence can be interpreted as an adjacency matrix. Furthermore, the vector $\mathbf{e}$, assumed to follow a zero-mean white Gaussian distribution, i.e., $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, accounts for the noise. This assumption is inspired by the findings in [6]. Therefore, the covariance of $\mathbf{x}$ can be derived as $\mathbb{E}[\mathbf{x}\mathbf{x}^\top] = \mathbb{E}[(\mathbf{I} - \mathbf{S})^{-1}\mathbf{e}\mathbf{e}^\top(\mathbf{I} - \mathbf{S})^{-\top}] = (\mathbf{I} - \mathbf{S})^{-1}(\mathbf{I} - \mathbf{S})^{-\top}$.

In case we consider multiple independent realizations of $\mathbf{e}$, which can be stacked in $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_T]$, we obtain multiple independent realizations of $\mathbf{x}$, grouped in $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$, as

$$\mathbf{X} = \mathbf{S}\mathbf{X} + \mathbf{E}. \tag{2}$$

Polynomial model: The SEM described in (1) (or equivalently (2)) essentially takes the form $\mathbf{x} = \mathbf{H}\mathbf{e}$, where $\mathbf{H} = (\mathbf{I} - \mathbf{S})^{-1}$. This representation can be viewed as a special case of a polynomial relationship because the matrix $\mathbf{H} = (\mathbf{I} - \mathbf{S})^{-1}$ can usually be expanded into a power series of $\mathbf{S}$. This concept forms the basis of the polynomial model, which considers

$$\mathbf{x} = h(\mathbf{S})\mathbf{e}. \tag{3}$$

In this model, $h(x)$ represents any polynomial of the form $h(x) = \sum_{l=0}^{L-1} h_l x^l$. We again assume that $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

III. Topology Identification for SEM

In this section, we focus on the SEM for undirected graphs, which we then later on will extend to any other known polynomial model or latent variable distribution. Note that we will use a hat notation to represent the optimization variable whereas we use a star to indicate the optimal solution.

A straightforward approach to estimate a symmetric hollow $\mathbf{S}$ from (2) is to minimize $\|\mathbf{X}-\hat{\mathbf{S}}\mathbf{X}\|_F^2$ using the constraints $\text{diag}(\hat{\mathbf{S}}) = \mathbf{0}$ and $\hat{\mathbf{S}} = \hat{\mathbf{S}}^\top$. However, it can be shown that this does not lead to the ground truth $\mathbf{S}$, even when the number of samples T grows to infinity and $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. To prove this, observe that for an infinite number of observations, the objective function can be defined as $f(\hat{\mathbf{S}}) = \mathbb{E}\{\|\mathbf{x} - \hat{\mathbf{S}}\mathbf{x}\|_2^2\} = \text{tr}((\mathbf{I}-\mathbf{S})^{-2}(\mathbf{I}-\hat{\mathbf{S}})^2)$. Given that $\hat{\mathbf{S}}$ is a symmetric matrix, for the optimal solution we should obtain $\frac{df}{d\hat{\mathbf{S}}} = \frac{\partial f}{\partial \hat{\mathbf{S}}} + \frac{\partial f}{\partial \hat{\mathbf{S}}^\top} - \text{Diag}(\frac{\partial f}{\partial \hat{\mathbf{S}}}) = \mathbf{\Lambda}$ [12], where $\mathbf{\Lambda}$ is a diagonal matrix because non-zero off-diagonal elements would allow gradient descent to further minimize f [1]. However, setting $\hat{\mathbf{S}}$ to the true $\mathbf{S}$, it can be proven that $\frac{\partial f}{\partial \hat{\mathbf{S}}}|_{\hat{\mathbf{S}}=\mathbf{S}} = -2(\mathbf{I}-\mathbf{S})^{-1}$. As a result, $\frac{df}{d\hat{\mathbf{S}}}|_{\hat{\mathbf{S}}=\mathbf{S}} = \mathbf{\Lambda}$ is only possible when $\mathbf{S}$ is a diagonal matrix, which is not possible.

Instead, we focus on a covariance matching approach, where we match the sample covariance matrix $\mathbf{C}_\mathbf{x} = \mathbf{X}\mathbf{X}^\top/T$ to the theoretical covariance matrix $\mathbf{\Sigma}_\mathbf{x} = \mathbb{E}\{\mathbf{x}\mathbf{x}^\top\}$ expected from the model. Under the conditions $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathbf{S} = \mathbf{S}^\top$, the latter is given by

$$\mathbf{\Sigma}_\mathbf{x} = (\mathbf{I}-\mathbf{S})^{-1}(\mathbf{I}-\mathbf{S})^{-\top} = (\mathbf{I}-\mathbf{S})^{-2}. \tag{4}$$

Introducing $\mathbf{H} = (\mathbf{I}-\mathbf{S})^{-1}$ as earlier, we thus obtain $\mathbf{\Sigma}_\mathbf{x} = \mathbf{H}^2$. Estimating $\mathbf{H}$ instead of $\mathbf{S}$, our covariance matching problem can be formulated as

$$\begin{aligned} \mathbf{H}^* = \arg\min_{\hat{\mathbf{H}}} \quad & \|\hat{\mathbf{H}}^2 - \mathbf{C}_\mathbf{x}\|_F^2 \\ \text{subject to} \quad & \text{diag}(\hat{\mathbf{H}}^{-1}) = \mathbf{1}, \\ & \hat{\mathbf{H}} = \hat{\mathbf{H}}^\top. \end{aligned} \tag{5}$$

This problem is hard to solve though, so we further tune this into a more manageable form.

Let the eigenvalue decomposition (EVD) of the symmetric matrix variable $\hat{\mathbf{H}}$ be given by $\hat{\mathbf{H}} = \hat{\mathbf{U}}\text{diag}(\hat{\boldsymbol{\lambda}})\hat{\mathbf{U}}^T$, which leads to[2] $\hat{\mathbf{H}}^2 = \hat{\mathbf{U}}\text{diag}(\hat{\boldsymbol{\lambda}}^2)\hat{\mathbf{U}}^\top$. This allows us to replace the matrix variable $\hat{\mathbf{H}}$ by two new variables $\hat{\mathbf{U}}$ and $\hat{\boldsymbol{\lambda}}$. Consider now also the EVD of $\mathbf{C}_\mathbf{x}$, which is given by $\mathbf{C}_\mathbf{x} = \mathbf{U}_\mathbf{x}\text{diag}(\boldsymbol{\lambda}_\mathbf{x})\mathbf{U}_\mathbf{x}^\top$. Then we can simplify problem (5) by setting $\hat{\mathbf{U}} = \mathbf{U}_\mathbf{x}$ and restricting the problem to the single vector variable $\hat{\boldsymbol{\lambda}}$. Problem (5) can then be approximated as

$$\begin{aligned} \boldsymbol{\lambda}^* = \arg\min_{\hat{\boldsymbol{\lambda}}} \quad & \|\hat{\boldsymbol{\lambda}}^2 - \boldsymbol{\lambda}_\mathbf{x}\|_2^2 \\ \text{subject to} \quad & \text{diag}(\mathbf{U}_\mathbf{x}\text{diag}(\hat{\boldsymbol{\lambda}}^{-1})\mathbf{U}_\mathbf{x}^\top) = \mathbf{1}. \end{aligned} \tag{6}$$

Since this constraint is still hard to handle, we further approximate the problem by setting the objective to zero and turning the constraint into an objective.

[1]We adopt the definition of matrix differentiation for structured matrices from [12]. Specifically, for a scalar function $g(\mathbf{A})$, we define $\frac{dg}{dA_{ij}} = \sum_{kl} \frac{\partial g}{\partial A_{kl}} \frac{\partial A_{kl}}{\partial A_{ij}}$.

[2]All powers of vectors should be considered as element-wise.

Setting the objective to zero means that $\hat{\boldsymbol{\lambda}}^2 = \boldsymbol{\lambda}_\mathbf{x}$. This however introduces a sign ambiguity, which we can interpret as our new variable. So introducing $\hat{\mathbf{q}} \in \{-1, 1\}^{N\times 1}$ we can change the variable $\hat{\boldsymbol{\lambda}}$ into the binary variable $\hat{\mathbf{q}}$ by setting $\hat{\boldsymbol{\lambda}} = \text{diag}(\hat{\mathbf{q}})\boldsymbol{\lambda}_\mathbf{x}^{1/2}$, where $(\cdot)^{1/2}$ represents the positive square root.

Now turning the constraint in problem (6) into an objective function, we obtain

$$\begin{aligned} & \|\,\text{diag}(\mathbf{U}_\mathbf{x}\text{diag}(\hat{\boldsymbol{\lambda}}^{-1})\mathbf{U}_\mathbf{x}^\top) - \mathbf{1}\|_2^2 \\ &= \|\text{diag}(\mathbf{U}_\mathbf{x}\text{diag}(\hat{\mathbf{q}}^{-1})\text{diag}(\boldsymbol{\lambda}_\mathbf{x}^{-1/2})\mathbf{U}_\mathbf{x}^\top) - \mathbf{1}\|_2^2 \\ &= \|\text{diag}(\mathbf{U}_\mathbf{x}\text{diag}(\hat{\mathbf{q}})\text{diag}(\boldsymbol{\lambda}_\mathbf{x}^{-1/2})\mathbf{U}_\mathbf{x}^\top) - \mathbf{1}\|_2^2 \\ &= \|(\mathbf{U}_\mathbf{x} \odot \mathbf{U}_\mathbf{x})\text{diag}(\boldsymbol{\lambda}_\mathbf{x}^{-1/2})\hat{\mathbf{q}} - \mathbf{1}\|_2^2, \end{aligned} \tag{7}$$

where $\odot$ is the element-wise (Hadamard) product and where we have used $\hat{\mathbf{q}} = \hat{\mathbf{q}}^{-1}$. Hence, the objective now becomes a simple quadratic function in the binary variables $\hat{\mathbf{q}}$. Defining $\mathbf{W} = (\mathbf{U}_\mathbf{x} \odot \mathbf{U}_\mathbf{x})\text{diag}(\boldsymbol{\lambda}_\mathbf{x}^{-1/2})$ our proposed problem can finally be stated as the following binary least squares problem also known as an unconstrained binary quadratic programming (UBQP) problem:

$$\textbf{(P1)} \qquad \mathbf{q}^* = \arg\min_{\hat{\mathbf{q}}\in\{-1,1\}^{N\times 1}} \|\mathbf{W}\hat{\mathbf{q}} - \mathbf{1}\|_2^2. \tag{8}$$

For this problem, we can state the following identifiabily theorem.

Theorem 1: Let the EVD of the true GSO $\mathbf{S}$ be given by $\mathbf{S} = \mathbf{U}\,\text{diag}(\boldsymbol{\lambda})\mathbf{U}^\top$. Further assume $\hat{\mathbf{p}}$ is a binary variable and consider the equation

$$(\mathbf{U} \odot \mathbf{U})|\mathbf{I} - \text{diag}(\boldsymbol{\lambda})|\hat{\mathbf{p}} - \mathbf{1} = \mathbf{0}. \tag{9}$$

If this equation only has one binary solution $\mathbf{p}^*$, then the estimator $\mathbf{S}^*$, obtained from the solution of problem **(P1)**, i.e., $\mathbf{S}^* = \mathbf{I} - \mathbf{U}_\mathbf{x}\,\text{diag}(\mathbf{q}^*)\,\text{diag}(\boldsymbol{\lambda}_\mathbf{x}^{-1/2})\mathbf{U}_\mathbf{x}^\top$, will converge to the true $\mathbf{S}$ when the number of observations T goes to infinity.

Our proof sketch starts with observing that at $T = \infty$ we have $\mathbf{C}_\mathbf{x} = \mathbf{\Sigma}_\mathbf{x} = (\mathbf{I}-\mathbf{S})^{-2}$. As a result, the EVD of $\mathbf{C}_\mathbf{x}$ then is $\mathbf{C}_\mathbf{x} = \mathbf{U}(\mathbf{I} - \text{diag}(\boldsymbol{\lambda}))^{-2}\mathbf{U}^\top$, and thus $\mathbf{U}_\mathbf{x} = \mathbf{U}$ and $\boldsymbol{\lambda}_\mathbf{x}^{-1/2} = |\mathbf{1} - \boldsymbol{\lambda}|$. Hence, saying that (9) has a unique binary solution $\mathbf{p}^*$ is the same as saying that (8) has a unique solution $\mathbf{q}^*$ at $T = \infty$ and these solutions are then also the same.

Although this theorem may seem evident, it can be considered as a broadening of the theorem mentioned in [13], where $\text{rank}(\mathbf{U}\odot\mathbf{U}) = N - 1$ is required. Under this condition, our theorem holds trivially. However, our theorem has the potential to handle cases where $\text{rank}(\mathbf{U} \odot \mathbf{U}) < N - 1$ and in our experiments, we will verify this.

There are many ways to solve an UBQP. Here we consider the traditional semi-definite relaxation approach and we refer the reader to [14] for more details.

IV. Extension to Polynomial Model

In this section, we extend the SEM approach to the polynomial model. If $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathbf{S} = \mathbf{S}^\top$, then according to (3) we have

$$\mathbf{\Sigma}_\mathbf{x} = h(\mathbf{S})h^\top(\mathbf{S}) = h^2(\mathbf{S}) \tag{10}$$

Similar as in Section III, we will then try to match $\mathbf{C}_\mathbf{x}$ to $h^2(\mathbf{S})$. To do that, we basically follow the same ideas as for the SEM and start by defining the EVD of the matrix variable $\hat{\mathbf{S}}$ as $\hat{\mathbf{S}} = \hat{\mathbf{U}}\,\text{diag}(\hat{\boldsymbol{\lambda}})\hat{\mathbf{U}}^\top$. Further decomposing $\mathbf{C}_\mathbf{x}$ as $\mathbf{C}_\mathbf{x} = \mathbf{U}_\mathbf{x}\,\text{diag}(\boldsymbol{\lambda}_\mathbf{x})\mathbf{U}_\mathbf{x}^\top$, we can set

$\hat{\mathbf{U}} = \mathbf{U_x}$. This allows us to match $h^2(\text{diag}(\hat{\boldsymbol{\lambda}}))$ to $\text{diag}(\boldsymbol{\lambda_x})$ under the hollow constraint $\text{diag}(\mathbf{U_x}\text{diag}(\hat{\boldsymbol{\lambda}})\mathbf{U_x}^\top) = (\mathbf{U_x} \odot \mathbf{U_x})\hat{\boldsymbol{\lambda}} = \mathbf{0}$. Switching again the objective and constraint, we finally obtain the problem

$$\begin{aligned} \boldsymbol{\lambda}^* = \arg\min_{\hat{\boldsymbol{\lambda}}} \quad & \|(\mathbf{U_x} \odot \mathbf{U_x})\hat{\boldsymbol{\lambda}}\|_2^2 \\ \text{subject to} \quad & h^2(\text{diag}(\hat{\boldsymbol{\lambda}})) - \text{diag}(\boldsymbol{\lambda_x}) = \mathbf{0}. \end{aligned} \tag{11}$$

The constraint basically represents a set of scalar polynomial constraints of the form $h^2(\hat{\lambda}_i) - \lambda_{i,\mathbf{x}} = 0$, $i = 1, 2, \ldots, N$, where $\hat{\lambda}_i$ ($\lambda_{i,\mathbf{x}}$) denotes the ith element of $\hat{\boldsymbol{\lambda}}$ ($\boldsymbol{\lambda_x}$). Denoting the roots of the ith scalar polynomial as $\mathcal{C}_i = \{c_i^1, c_i^2, \ldots, c_i^{p_i}\}$ we can replace $h^2(\hat{\lambda}_i) - \lambda_{i,\mathbf{x}} = 0$ by $\hat{\lambda}_i \in \mathcal{C}_i$. Our proposed problem can finally be stated as

$$\begin{aligned} \textbf{(P2)} \qquad \boldsymbol{\lambda}^* = \arg\min_{\hat{\boldsymbol{\lambda}}} \quad & \|(\mathbf{U_x} \odot \mathbf{U_x})\hat{\boldsymbol{\lambda}}\|_2^2 \\ \text{subject to } \hat{\lambda}_i \in \mathcal{C}_i, \; & i = 1, 2, \ldots, N. \end{aligned} \tag{12}$$

Let us explore how **(P2)** specializes to **(P1)**. If we set $h(\mathbf{S}) = (\mathbf{I} - \mathbf{S})^{-1}$ in **(P2)**, then all solutions to the equation $h^2(\text{diag}(\hat{\boldsymbol{\lambda}})) - \text{diag}(\boldsymbol{\lambda_x}) = \mathbf{0}$ can be expressed as $\hat{\boldsymbol{\lambda}} = \text{diag}(\boldsymbol{\lambda_x}^{-1/2})\hat{\mathbf{q}} - \mathbf{1}$, where $\hat{\mathbf{q}} \in \{-1, 1\}^{N \times 1}$. Therefore, the objective of **(P2)** can be rewritten as $(\mathbf{U_x} \odot \mathbf{U_x})(\text{diag}(\boldsymbol{\lambda_x}^{-1/2})\hat{\mathbf{q}} - \mathbf{1})$, which, due to the property $(\mathbf{U_x} \odot \mathbf{U_x})\mathbf{1} = \mathbf{1}$, becomes identical to the problem defined in **(P1)**. Thus, **(P2)** can be viewed as an extension of **(P1)**.

Similar to the semi-definite relaxation approach used for the SEM problem, this problem can also be solved using convex relaxation approaches [15]. Additionally, solvers like Gurobi [16] that support integer programming can also be used to find a global minimum.

V. Extension to General Distribution

In this section, we go back to the regular SEM and explore extending the distribution of the latent variable $\mathbf{e}$. More specifically, consider $\mathbf{e}$ to be normally distributed as $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma_e})$, where $\boldsymbol{\Sigma_e}$ is known.

Estimating $\mathbf{S}$ then again boils down to estimating $\mathbf{H} = (\mathbf{I} - \mathbf{S})^{-1}$. At first sight, one could exploit the fact that $\boldsymbol{\Sigma_x} = \mathbf{H}\boldsymbol{\Sigma_e}\mathbf{H}^\top$ and match $\hat{\mathbf{H}}\boldsymbol{\Sigma_e}\hat{\mathbf{H}}^\top$ with $\mathbf{C_x}$. Solving this matching problem is challenging though. As an alternative, observe that $(\mathbf{H}\boldsymbol{\Sigma_e})^2 = \mathbf{H}\boldsymbol{\Sigma_e}\mathbf{H}\boldsymbol{\Sigma_e} = \boldsymbol{\Sigma_x}\boldsymbol{\Sigma_e}$. This allows us to match $(\hat{\mathbf{H}}\boldsymbol{\Sigma_e})^2$ with $\mathbf{C_x}\boldsymbol{\Sigma_e}$ which is similar to (5), where we matched $\hat{\mathbf{H}}^2$ with $\mathbf{C_x}$. As a result, we follow again the same steps.

First, we introduce two new variables $\hat{\mathbf{U}}$ and $\hat{\boldsymbol{\lambda}}$ by considering the EVD of $\hat{\mathbf{H}}\boldsymbol{\Sigma_e}$, i.e., $\hat{\mathbf{H}}\boldsymbol{\Sigma_e} = \hat{\mathbf{U}}\,\text{diag}(\hat{\boldsymbol{\lambda}})\hat{\mathbf{U}}^{-1}$. This obviously leads to $(\hat{\mathbf{H}}\boldsymbol{\Sigma_e})^2 = \hat{\mathbf{U}}\,\text{diag}(\hat{\boldsymbol{\lambda}})^2\hat{\mathbf{U}}^{-1}$. Computing the EVD of $\mathbf{C_x}\boldsymbol{\Sigma_e}$, we obtain[3] $\mathbf{C_x}\boldsymbol{\Sigma_e} = \mathbf{U_{xe}}\,\text{diag}(\boldsymbol{\lambda_{xe}})\mathbf{U_{xe}}^{-1}$. Setting now $\hat{\mathbf{U}} = \mathbf{U_{xe}}$ and replacing the matching problem by the constraint $\hat{\boldsymbol{\lambda}}^2 = \boldsymbol{\lambda_{xe}}$ introduces once again a sign ambiguity. More specifically, we can change the variable $\hat{\boldsymbol{\lambda}}$ by the binary variable $\hat{\mathbf{q}} \in \{-1, 1\}^{N \times 1}$ using $\hat{\boldsymbol{\lambda}} = \text{diag}(\hat{\mathbf{q}})\boldsymbol{\lambda_{xe}}^{1/2}$. Overall, this allows us to write $\hat{\mathbf{H}}$ as a function of $\hat{\mathbf{q}}$ through

$$\hat{\mathbf{H}} = \mathbf{U_{xe}}\,\text{diag}(\hat{\mathbf{q}})\text{diag}(\boldsymbol{\lambda_{xe}}^{1/2})\mathbf{U_{xe}}^{-1}\boldsymbol{\Sigma_e}^{-1}. \tag{13}$$

The inverse of $\hat{\mathbf{H}}$ is then given by

$$\hat{\mathbf{H}}^{-1} = \boldsymbol{\Sigma_e}\mathbf{U_{xe}}\text{diag}(\boldsymbol{\lambda_{xe}}^{-1/2})\,\text{diag}(\hat{\mathbf{q}})\mathbf{U_{xe}}^{-1} \tag{14}$$

[3]Note that we use the notation $\mathbf{U}$ primarily to align with the previous notation and it does not imply that $\mathbf{U}$ is unitary

and the diagonal of $\hat{\mathbf{H}}^{-1}$ is

$$\text{diag}(\hat{\mathbf{H}}^{-1}) = [(\boldsymbol{\Sigma_e}\mathbf{U_{xe}}) \odot \mathbf{U_{xe}}^{-\top}]\text{diag}(\boldsymbol{\lambda_{xe}}^{-1/2})\hat{\mathbf{q}}. \tag{15}$$

Finally, defining $\mathbf{W} = [(\boldsymbol{\Sigma_e}\mathbf{U_{xe}}) \odot \mathbf{U_{xe}}^{-\top}]\text{diag}(\boldsymbol{\lambda_{xe}}^{-1/2})$, the optimization problem simplifies to:

$$\textbf{(P3)} \qquad \min_{\hat{\mathbf{q}} \in \{1,1\}^{N \times 1}} \|\mathbf{W}\hat{\mathbf{q}} - \mathbf{1}\|_2^2, \tag{16}$$

which is again an UBQP that can be solved using semi-definite relaxation.

Comparing **(P1)** and **(P3)**, their formulations are almost identical. Furthermore, if we assume $\boldsymbol{\Sigma_e} = \mathbf{I}$, then $\mathbf{U_{xe}}$ in **(P3)** reduces to $\mathbf{U_x}$, and $\mathbf{U_{xe}}^{-\top}$ also reduces to $\mathbf{U_x}$. Moreover, $\boldsymbol{\lambda_{xe}}^{-1/2}$ reduces to $\boldsymbol{\lambda_x}^{-1/2}$. This extension of the SEM problem is quite elegant, as it scarcely alters the structure of the problem, and the complexity of solving the optimization problem remains the same.

VI. Experiments

Here we consider some experiments using simulated and real graphs. For the simulated graphs, we generate 100 realizations. We evaluate the normalized squared error, defined as $\text{NSE}(\mathbf{S}, \mathbf{S}^*) = \|\mathbf{S}^* - \mathbf{S}\|_F^2 / \|\mathbf{S}\|_F^2$, averaged over these 100 graphs across various sample sizes T ranging from 10^2 up to 10^6.

Comparison with other methods:

Here, we compare our approach (refered to as CovMatch) with SpecTemp [13] and with a trivial signal matching approach (SigMatch) [4] based on minimizing $\|\mathbf{X} - \hat{\mathbf{S}}\mathbf{X}\|_F^2$. Due to the sign constraints of SpecTemp, all graphs are assigned positive weights ranging from 0.1 to 2. At the same time, we adopt the simplest assumption that $\boldsymbol{\Sigma_e} = \mathbf{I}$. In the first scenario (labelled as simple), we deliberately generate graphs where $\text{rank}(\mathbf{U} \odot \mathbf{U}) = N - 1$. In the second scenario (labelled as hard), we only generate graphs with $\text{rank}(\mathbf{U} \odot \mathbf{U}) < N - 1$ to check the robustness of our method. In both scenarios, the number of nodes and edges are set to 20. For each scenario, we calculate the average NSE. Note that a singular value less than 5×10^{-4} is considered as a rank loss here.

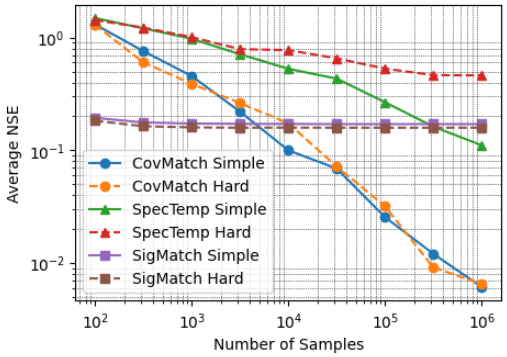

Fig. 1: Average NSE for different samples.

As shown in Fig. 1, it is evident that SpecTemp often fails due to a loss of rank. Conversely, our method, CovMatch, continues to perform well. This highlights the robustness and reliability of CovMatch. Further, the SigMatch approach never converges to the correct result, but it performs better than others with less observations.

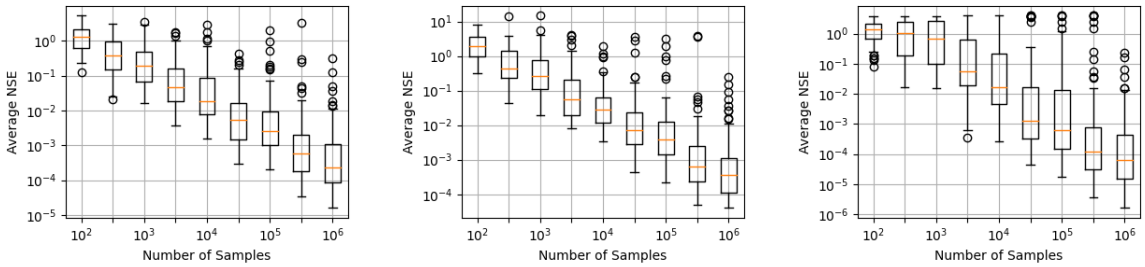

(a) SEM sparse graph (b) SEM fully connected graph (c) Polynomial model sparse graph

Fig. 2: Distribution of NSE across different graph configurations and methods. (a) SEM sparse graph, (b) SEM fully connected graph, and (c) polynomial model. As $T \to \infty$, the NSE values are all smaller than 2×10^{-5}.

This is because both SpecTemp and CovMatch highly rely on an accurate sample covariance, which requires many samples.

Besides the advantages mentioned above, our method can also handle negative edge weights and some very challenging graphs. For the next experiment, the edge weights of the 100 graph realizations are chosen within the range $[-2, -0.1] \cup [0.1, 2]$. For the covariance matrix of the latent variables $\mathbf{\Sigma_e}$, we begin by generating a random $N \times N$ matrix, $\mathbf{\Sigma}_{\text{half}}$, with each element uniformly distributed between $[-1, 1]$. The covariance matrix is then formed as $\mathbf{\Sigma_e} = \mathbf{\Sigma}_{\text{half}}\mathbf{\Sigma}_{\text{half}}^{\top}$, ensuring it is symmetric and positive semi-definite. For the SEM, we generate $\mathbf{\Sigma_e}$ according to this method, while for the polynomial model, we still adopt $\mathbf{\Sigma_e} = \mathbf{I}$.

SEM for sparse graphs: In this scenario, we employ a graph with 20 edges and 20 nodes, and intentionally create a challenging condition by generating graphs with a rank of $N - 3$. This condition may cause some methods to fail due to the presence of repeated eigenvalues in $\mathbf{C_x}$ [11]. However, our method remains effective under these constraints.

SEM for fully connected graphs: For the second configuration, we test our method on a fully connected graph and we ignore the rank constraint. Methods that rely on sparsity often fail in this scenario, but our approach continues to perform well.

Polynomial model: Here we consider graphs with 20 nodes and 40 edges while ignoring the rank constraint and we apply the polynomial model. Additionally, we define the polynomial function $h(x)$ as a third-order polynomial $h(x) = \sum_{i=0}^{3} h_i x^i$, where each coefficient h_i is randomly chosen from the interval $[-1, 1]$.

As demonstrated in Fig. 2, our methodology has been notably successful not only on two particularly challenging graph configurations for a SEM but also on nontrivial polynomial models. As the sample size increases, the average NSE generally decreases to significantly low levels. In ideal conditions, with an infinite sample size, our errors can approach zero, showcasing the robustness and effectiveness of our approach. These results also provide indirect confirmation of the correctness of Theorem 1.

Real data: We also compare our approach with network deconvolution (referred to as NetDeconv) [17], which similarly involves estimating $\mathbf{S}$ from $\mathbf{H}$. In this experiment, each node within the network corresponds to an amino acid residue, and the edges denote mutual information, reflecting co-variation among residues across multiple sequence alignments that include 2,000 to 72,000 sequences. Our objective is to deduce structural constraints among amino acid pairs to aid in predicting protein structures.

We employ a relatively straightforward polynomial $h(x) = \frac{1}{80}(x^3 + 2x^2 + 4x)$, whereas NetDeconv approximates $h(x) = \frac{x}{1-x}$. Terms such as "1wvn" shown in Figure 3 represent different protein labels. The figure illustrates that in various cases, our results outperform those of NetDeconv.

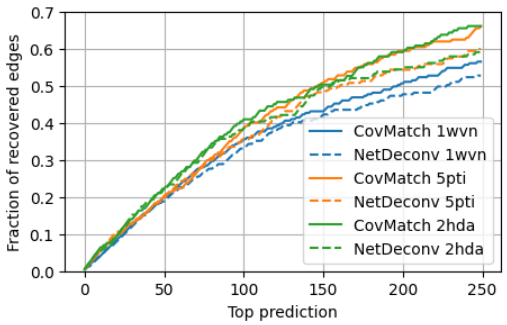

Fig. 3: Real contact edge recovery as a function of the number of edges considered.

VII. Conclusions

In this paper, we have introduced a promising topology identification methodology based on covariance matching, which is fundamentally based on reproducing the theoretical covariance model from the sample covariance matrix. This approach has significant potential due to its foundation in regenerating observable data characteristics. Focusing on the SEM as our primary area of study, we have described how we can simplify the problem into a UBQP which can be solved by semi-definite relaxation. For this UBQP we also provide a convergence proof to the true graph. Furthermore, we assert that our method can be extended to more complex scenarios, such as towards the polynomial model framework, and even to any known distribution of latent variables. We substantiate the efficacy and correctness of our approach through extensive experimental validation, demonstrating its robustness across a variety of settings.

References

[1] Jiaying Liu, Feng Xia, Lei Wang, Bo Xu, Xiangjie Kong, Hanghang Tong, and Irwin King, "Shifu2: A network representation learning based model for advisor-advisee relationship mining," *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 4, pp. 1763–1777, 2019.

[2] AR Mclntosh and Francisco Gonzalez-Lima, "Structural equation modeling and its application to network analysis in functional brain imaging," *Human brain mapping*, vol. 2, no. 1-2, pp. 2–22, 1994.

[3] Judea Pearl, "Graphs, causality, and structural equation models," *Sociological Methods & Research*, vol. 27, no. 2, pp. 226–284, 1998.

[4] Alberto Natali, Elvin Isufi, Mario Coutino, and Geert Leus, "Online graph learning from time-varying structural equation models," in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 1579–1585.

[5] Juan Andrés Bazerque, Brian Baingana, and Georgios B Giannakis, "Identifiability of sparse structural equation models for directed and cyclic networks," in *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, 2013, pp. 839–842.

[6] J. Peters and P. Bühlmann, "Identifiability of Gaussian structural equation models with equal error variances," *Biometrika*, vol. 101, no. 1, pp. 219–228, 11 2013.

[7] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing, "Dags with no tears: Continuous optimization for structure learning," *Advances in neural information processing systems*, vol. 31, 2018.

[8] Seyed Saman Saboksayr, Gonzalo Mateos, and Mariano Tepper, "Block successive convex approximation for concomitant linear dag estimation," in *2024 IEEE 13rd Sensor Array and Multichannel Signal Processing Workshop (SAM)*. IEEE, 2024, pp. 1–5.

[9] Gonzalo Mateos, Santiago Segarra, Antonio G Marques, and Alejandro Ribeiro, "Connecting the dots: Identifying network structure via graph signal processing," *IEEE Signal Processing Magazine*, vol. 36, no. 3, pp. 16–43, 2019.

[10] Alberto Natali, Elvin Isufi, Mario Coutino, and Geert Leus, "Learning time-varying graphs from online data," *IEEE Open Journal of Signal Processing*, vol. 3, pp. 212–228, 2022.

[11] Andrei Buciulea, Jiaxi Ying, Antonio G. Marques, and Daniel P. Palomar, "Polynomial graphical lasso: Learning edges from gaussian graph-stationary signals," 2024.

[12] Kaare Brandt Petersen, Michael Syskind Pedersen, et al., "The matrix cookbook," *Technical University of Denmark*, vol. 7, no. 15, pp. 510, 2008.

[13] Santiago Segarra, Antonio G Marques, Gonzalo Mateos, and Alejandro Ribeiro, "Network topology inference from spectral templates," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 3, pp. 467–483, 2017.

[14] Zhi-Quan Luo, Wing-Kin Ma, Anthony Man-Cho So, Yinyu Ye, and Shuzhong Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, 2010.

[15] Christoph Buchheim and Angelika Wiegele, "Semidefinite relaxations for non-convex quadratic mixed-integer programming," *Mathematical Programming*, vol. 141, pp. 435–452, 2013.

[16] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2024.

[17] Soheil Feizi, Daniel Marbach, Muriel Médard, and Manolis Kellis, "Network deconvolution as a general method to distinguish direct dependencies in networks," *Nature biotechnology*, vol. 31, no. 8, pp. 726–733, 2013.

Unclonable Encryption with Continuous Variables

Arpan Akash Ray and Boris Škorić

Technische Universiteit Eindhoven, The Netherlands

February 28, 2025

Abstract

We propose the first continuous-variable (CV) unclonable encryption scheme, extending the paradigm of quantum encryption of classical messages (QECM) to CV systems. In our construction, a classical message is first encrypted classically and then encoded using an error-correcting code. Each bit of the codeword is mapped to a CV mode by creating a coherent state which is squeezed in the q or p quadrature direction, with a small displacement that encodes the bit. The squeezing directions are part of the encryption key. We prove unclonability in the framework introduced by Broadbent and Lord, via a reduction of the cloning game to a CV monogamy-of-entanglement game.

1 Introduction

The marriage of quantum information theory with cryptography has given rise to a wide array of protocols that exploit uniquely quantum phenomena, most notably the no-cloning principle [1, 2, 3], to achieve security properties that are unattainable in a classical setting. One such concept, *unclonable encryption*, harnesses the indivisibility of quantum states to prevent an adversary from copying an encrypted message.

The term Unclonable Encryption (UE) first appeared in 2003 in a paper by Gottesman [4]. Alice encrypts a classical message into a quantum state. A security definition was introduced that essentially states "*If Bob decides that his decryption is valid, then Eve, given the key, has only negligible information about the plaintext.*" The security definition was formulated in terms of the trace distance between encryptions of different messages. In the same framework, Leermakers and Škorić devised an UE scheme with key recycling [5]. Broadbent and Lord [6] introduced a modified security definition, based on a cloning game, and constructed UE in the random oracle model. Several further UE schemes were introduced in [7, 8, 9], and results on the feasibility and limitations of UE were given in [10, 11].

Until now, UE has been exclusively studied in discrete-variable (DV) quantum systems. However, in the field of Quantum Key Distribution (QKD), continuous-variable (CV) quantum systems have emerged as an attractive alternative to DV [12, 13, 14, 15, 16, 17, 18, 19, 20, 21]. CV does not need expensive single-photon detectors, and has the advantage that low-loss telecom wavelengths (1310nm, 1550nm) can be used, making it possible to capitalize on several decades of experience in coherent optical communication technology. Beyond QKD, the practical advantages hold more generally for other quantum information processing applications, and this has fueled substantial interest in translating DV-based cryptographic ideas into the CV domain.

In this paper, we propose the first Unclonable Encryption scheme that works with Continuous-Variable states. We provide a security proof in the UE framework of [6]. It turns out that bringing UE from discrete to continuous variables has a number of nontrivial aspects. On the construction side, the parameters of the scheme need to be tuned such that there both decryptability and unclonability are satisfied. On the proof-technical side, we introduce a number of 'game hops' in order to connect the *cloning game*, which features in the UE security definition, to the *CV monogamy-of-entanglement game*, for which an upper bound on the winning probability has recently been proven [22]. Furthermore it is necessary to slightly modify the definition of an

encryption, in order to allow for a small probability of decryption failure, caused by the probability tails of the CV measurements.

Beyond the immediate theoretical interest, a CV-based UE scheme has potential advantages for future practical quantum networks, especially where CV platforms are more readily integrated with existing optical infrastructure.

The outline of the paper is as follows. In Section 2 we briefly review CV formalism and important definitions and results from the literature. In Section 3 we present our protocol and verify that it satisfies the definition of a quantum encryption. Section 4 contains the unclonability proof. The main result is stated in Theorem 4.4. In Section 5 we summarize and discuss future work.

2 Preliminaries

2.1 Notation

We use standard bra-ket notation for quantum states. Hilbert spaces are written as $\mathcal{H}$ with a subscript. E.g. we write the Hilbert space of a single CV mode as $\mathcal{H}_1$. The notation $\mathcal{D}(\mathcal{H})$ stands for the set of density operators on $\mathcal{H}$. The Hamming weight of a string s is written as $|s|$.

2.2 Continuous Variables; Gaussian states

A 'mode' of the electromagnetic vector potential represents a plane wave solution of the vacuum Maxwell equations at a certain frequency, wave vector and polarisation. Associated with each mode there is a creation operator $\hat{a}^\dagger$ and annihilation operator $\hat{a}$. The linear combinations $\hat{q} = \frac{\hat{a}+\hat{a}^\dagger}{\sqrt{2}}$ and $\hat{p} = \frac{\hat{a}-\hat{a}^\dagger}{i\sqrt{2}}$ are easy-to-observe quantities called *quadratures*, and they behave as the position and momentum operator of a harmonic oscillator.

The *Gaussian states* are a special class of CV states; their Wigner function (quasi density function on the phase space) [23] is a Gaussian function of the quadrature variables. An N-mode Gaussian state is fully characterized by a displacement vector $d \in \mathbb{R}^{2N}$ and $2N \times 2N$ covariance matrix Γ. The corresponding Wigner function is $\frac{1}{\pi^N\sqrt{\det\Gamma}} \exp -(x-d)^T\Gamma^{-1}(x-d)$, where x stands for the vector $(q_1, p_1, \ldots, q_N, p_N)^T$. The class of Gaussian states contains important states like the vacuum, thermal states, coherent states, squeezed coherent states and EPR states (two-mode squeezed vacuum). A coherent state that is squeezed in the q-direction has covariance matrix $\left(\begin{smallmatrix} e^{-\zeta} & 0 \\ 0 & e^{\zeta} \end{smallmatrix}\right)$, where $\zeta \geq 0$ is the squeezing parameter. For the p-direction it is $\left(\begin{smallmatrix} e^{\zeta} & 0 \\ 0 & e^{-\zeta} \end{smallmatrix}\right)$. All intermediate directions are possible, but will not be used in this paper.

An EPR state is obtained by mixing a q-squeezed vacuum with a p-squeezed vacuum using a 50/50 beam splitter. The resulting two-mode squeezed (TMS) state has zero displacement, and its covariance matrix is $\left(\begin{smallmatrix} \mathbb{I}\cosh\zeta & \sigma_z \sinh\zeta \\ \sigma_z \sinh\zeta & \mathbb{I}\cosh\zeta \end{smallmatrix}\right)$, where $\mathbb{I}$ is the 2×2 identity matrix and $\sigma_z = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$. In the entanglement based version of our scheme we will make use of a state that resembles a displaced EPR state, with for instance displacement $d = (\alpha, 0, \alpha, 0)^T$. When the q-quadrature is measured on one side, the probability density for measurement outcome x_A is a normal distribution with mean α and variance $\frac{1}{2}\cosh\zeta$ (measuring a single quadrature is called a *homodyne* measurement). The measurement projects the state on the other side to a squeezed coherent state with displacement $(\alpha + [x_A - \alpha]\tanh\zeta, 0)^T$ and covariance matrix $\left(\begin{smallmatrix} \frac{1}{\cosh\zeta} & 0 \\ 0 & \cosh\zeta \end{smallmatrix}\right)$. If the x_A is not known, the displacement $\alpha + [x_A - \alpha]\tanh\zeta$ is a stochastic variable following a Gaussian distribution with mean α and variance $\frac{1}{2}\cosh\zeta\tanh^2\zeta$.

For a comprehensive review of CV quantum information we refer to [23].

2.3 Definitions and useful lemmas

In Section 3 we will introduce a scheme that encrypts a classical message into a quantum ciphertext (cipherstate), using a classical key. For the formal description of such a scheme we follow the definition given in [6], with a small modification: we allow for a small probability of failure in the decryption.

Definition 2.1. *Let $\lambda \in \mathbb{N}^+$ be a security parameter. Let $\mathcal{M}(\lambda)$ be the (classical) plaintext space. Let $\mathcal{K}(\lambda)$ be the (classical) key space. A* **quantum encryption of classical messages (QECM)** *scheme is a triplet of efficient quantum circuits $(Key_\lambda, Enc_\lambda, Dec_\lambda)$ implementing CPTP maps of the form*

$$Key_\lambda : \mathcal{D}(\mathcal{C}) \to \mathcal{D}(\mathcal{H}_{\mathcal{K}(\lambda)}) \tag{1}$$
$$Enc_\lambda : \mathcal{D}(\mathcal{H}_{\mathcal{K}(\lambda)} \otimes \mathcal{H}_{\mathcal{M}(\lambda)}) \to \mathcal{D}(\mathcal{H}_{T,\lambda}) \tag{2}$$
$$Dec_\lambda : \mathcal{D}(\mathcal{H}_{\mathcal{K}(\lambda)} \otimes \mathcal{H}_{T,\lambda}) \to \mathcal{D}(\mathcal{H}_{\mathcal{M}(\lambda)}) \tag{3}$$

where $\mathcal{H}_{\mathcal{K}(\lambda)}$ is the Hilbert space of the key, $\mathcal{H}_{\mathcal{M}(\lambda)}$ is the Hilbert space of the plaintext, and $\mathcal{H}_{T,\lambda}$ is the cipherstate space. Let E_k denote the CPTP map $\rho \mapsto Enc_\lambda(|k\rangle\langle k| \otimes \rho)$, and let D_k stand for the CPTP map $\rho \mapsto Dec_\lambda(|k\rangle\langle k| \otimes \rho)$. For all $k \in \mathcal{K}(\lambda)$, $m \in \mathcal{M}(\lambda)$, it must hold that

$$\mathrm{Tr}\left[|k\rangle\langle k| Key_\lambda(1)\right] > 0 \quad \Longrightarrow \quad \mathrm{Tr}\left[|m\rangle\langle m| \, \mathrm{D}_k \circ \mathrm{E}_k(|m\rangle\langle m|)\right] \geq 1 - \varepsilon_{\mathrm{DF}} \tag{4}$$

where $\varepsilon_{\mathrm{DF}}$ is the tolerated probability of decryption failure.

Unclonability of a QECM scheme is defined via a *cloning game*. A challenger prepares a cipherstate and gives it to Alice. Alice splits the cipherstate into two pieces; one piece goes to Bob, one to Charly. Then Bob and Charly receive the decryption key and must *both* produce the correct plaintext, without being allowed to communicate. A QECM scheme is considered to be secure against cloning if the three players ABC, acting together, have an exponentially small probability of winning the game.

We follow the definitions of [6] for the cloning game and the security against cloning.

Definition 2.2 (Cloning attack)**.** *Let S be a QECM scheme, with Hilbert spaces $\mathcal{H}_{\mathcal{K}(\lambda)}$, $\mathcal{H}_{\mathcal{M}(\lambda)}$, $\mathcal{H}_{T,\lambda}$ as given in Def. 2.1. Let $\mathcal{H}_B$ and $\mathcal{H}_C$ be Hilbert spaces of arbitrary dimension. A cloning attack against S is a triplet of efficient quantum circuits $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ implementing CPTP maps of the form*

$$\mathcal{A} : \mathcal{D}(\mathcal{H}_{T,\lambda}) \to \mathcal{D}(\mathcal{H}_B \otimes \mathcal{H}_C), \tag{5}$$
$$\mathcal{B} : \mathcal{D}(\mathcal{H}_{\mathcal{K}(\lambda)} \otimes \mathcal{H}_B) \to \mathcal{D}(\mathcal{H}_{\mathcal{M}(\lambda)}), \text{ and} \tag{6}$$
$$\mathcal{C} : \mathcal{D}(\mathcal{H}_{\mathcal{K}(\lambda)} \otimes \mathcal{H}_C) \to \mathcal{D}(\mathcal{H}_{\mathcal{M}(\lambda)}). \tag{7}$$

Definition 2.3. *Consider the cloning attack according to Def. 2.2. Let B_k stand for the CPTP map $\rho \mapsto \mathcal{B}(|k\rangle\langle k| \otimes \rho)$, and analogously C_k. Let $\mathcal{M}(\lambda) = \{0,1\}^n$. A QECM scheme S is $\tau(\lambda)$-*__uncloneable secure__ *if for all cloning attacks $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ against S, there exists a negligible function η such that*

$$\mathbb{E}_{m,k} \, \mathrm{Tr}\left(\left(|m\rangle\langle m| \otimes |m\rangle\langle m|\right) \left(B_k \otimes C_k\right) \circ \mathcal{A} \circ E_k\left(|m\rangle\langle m|\right) \right) \leq 2^{-n+\tau(\lambda)} + \eta(\lambda). \tag{8}$$

Here the expectation $\mathbb{E}$ is over uniform m, and k distributed according to $\Pr[K = k] = \mathrm{Tr}\left[|k\rangle\langle k| Key_\lambda(1)\right]$. If S is 0-uncloneable secure, we simply say that it is uncloneable secure.

Our security proof will make use of recent results on *entanglement monogamy games* for CV systems [22]. Such a game is played between Alice on one side and Bob and Charlie on the other. Bob and Charlie prepare a tripartite state ρ_{ABC} of their choice. Then Alice does an unpredictable measurement, and Bob and Charlie have to show that they are both sufficiently entangled with Alice to guess Alice's outcome to some degree of accuracy. Below we present the *coset monogamy game on* $\mathbb{R}^N$ as introduced in Section 4 of [22]. In the original version, the measurement bases are continuous variable coset states of $\mathbb{R}^N$, given by $|P_{q,p}\rangle = \bigotimes_{i=1}^{N} \begin{cases} |q = q_i\rangle, & i \notin I, \\ |p = p_i\rangle, & i \in I. \end{cases}$ Here, I is a subset interval of N. We note that this is equivalent to performing a homodyne measurement in the q direction, if $i \in I$, otherwise in the p direction. We refer to this game as the 'partial' CV monogamy of entanglement game, since it requires Bob and Charlie to produce guesses for only one quadrature direction instead of two. In Section 4.4 we introduce a 'full' version of the game.

Game 2.4: CV Partial Monogamy of Entanglement Game

1. **Initial state preparation.**
   Bob and Charlie prepare an M-mode state ρ_{ABC} across three registers: one for Alice, one for Bob, and one for Charlie. After preparation, they are no longer allowed to communicate.

2. **Alice's measurement choice and outcomes.**
   Alice chooses quadrature directions $(\theta_i)_{i=1}^{M}$, $\theta_i \in \{0, \frac{\pi}{2}\}$, such that $|i : \ \theta_i = 0| = M/2$, i.e. each direction is chosen exactly $M/2$ times. She then does homodyne detection of mode i in direction θ_i, getting outcomes that we denote as

$$q_i \quad \text{if } \theta_i = 0, \qquad p_i \quad \text{if } \theta_i = \tfrac{\pi}{2}.$$

3. **Announcement and responses.**
   Alice announces $(\theta_i)_{i=1}^{M}$. Bob outputs $q_B \in \mathbb{R}^{M/2}$ containing his guesses for Alice's q-values. Charlie outputs $p_C \in \mathbb{R}^{M/2}$ containing his guesses for Alice's p-values.

4. **Winning condition.**
   Bob and Charlie win the game if

$$\| q - q_B \|_\infty < \delta \quad \text{and} \quad \| p - p_C \|_\infty < \varepsilon. \tag{9}$$

The ∞-norm in the winning condition means that all guesses $(q_B)_i$ and $(p_C)_i$ must be close to Alice's values.
As shown in [22], this game can viewed as an *abelian coset measure* monogamy game, where Bob and Charlie attempt to guess the measurement outcomes in each quadrature within error neighborhoods $(-\delta, \delta)^N$ and $(-\varepsilon, \varepsilon)^N$. The following upper bound was obtained on the winning probability.

Lemma 2.5. *(Theorem 4.1 in [22]). In the CV Partial Monogamy of Entanglement Game (2.4), the winning probability w is upper bounded as*

$$w \ \leq \ \frac{1}{\binom{N}{N/2}} \sum_{k=0}^{N/2} \binom{N/2}{k}^2 \left(2\sqrt{\delta\,\varepsilon}\right)^k \ \leq \ \sqrt{e}\left(\tfrac{1}{2} + \sqrt{\delta\,\varepsilon}\right)^{\frac{N}{2}}. \tag{10}$$

3 Protocol Description

We propose a QECM scheme that makes use of squeezed coherent states. The scheme encrypts a message $m \in \{0,1\}^n$ to an N-mode cipherstate. First we apply a symmetric classical encryption scheme; this step ensures message confidentiality independent of the unclonability. We do not specify which symmetric cipher is used. We denote the encryption and decryption operations as $\texttt{Enc}_{\text{base}}(\cdot,\cdot)$ and $\texttt{Dec}_{\text{base}}(\cdot,\cdot)$.

Then the classical ciphertext gets encoded into an N-bit codeword c. The error correcting code is chosen such that it that can correct t bit errors. The encoding and decoding algorithms are denoted as `Encode`, `Decode`.

We encode a bit value c_i into a CV mode by displacing a squeezed vacuum state over a distance $(-1)^{c_i}\alpha$, where $\alpha > 0$ is a constant. The displacement is in the 'narrow' direction of the state. The underlying idea is that the squeezing direction is part of the encryption key; the attacker $\mathcal{A}$, who does not know this direction, has trouble determining c and hence cannot create good clones. This encoding is similar to conjugate coding for qubits [24].

For proof-technical reasons, the squeezing directions are not chosen uniformly at random. We impose the constraint that exactly half the modes are squeezed in the q-direction, and one half in the p-direction. Hence the corresponding key space is not $\{0,1\}^N$ but rather the set $\mathcal{L} = \{1, \ldots, \log_2 \binom{N}{N/2}\}$ of labels which uniquely enumerate the strings with Hamming weight $N/2$.

The QECM algorithms (Key, Enc, Dec) are given below.

Algorithm 1 Key Generation (Key)

Input: $z(\lambda)$, $N(\lambda)$, $r(\lambda)$, $\alpha(\lambda)$.
Output: a symmetric key $s \in \{0,1\}^z$, a binary string $\phi \in \{0,1\}^N$ and a real vector $k \in \mathbb{R}^N$.

1: Sample s uniformly from $\{0,1\}^z$.
2: Sample label ℓ uniformly from $\{1, \dots, \log_2 \binom{N}{N/2}\}$. Convert ℓ to a string $\phi \in \{0,1\}^N$ with Hamming weight $N/2$.
3: **for** $i = 1$ to N **do**
4: Sample k_i from the normal distribution $\mathcal{N}(0, \frac{1}{2}\cosh r \tanh^2 r)$ truncated to the interval $(-\alpha \tanh r, \alpha \tanh r)$.
5: **end for**
6: $k = (k_1, \dots, k_N)$.
7: Output (s, ϕ, k).

Note that the pdf of k_i is proportional to $\exp(-\frac{k_i^2}{\cosh r \tanh^2 r})$, i.e. Gaussian form, but the support $k_i \in (-\alpha \tanh r, \alpha \tanh r)$ is quite narrow compared to the width of the Gaussian.

Algorithm 2 Encryption (Enc)

Input: Key $(s, \phi, k) \in \{0,1\}^z \times \{0,1\}^N \times \mathbb{R}^{\mathbb{N}}$; message $m \in \{0,1\}^n$; parameters $\alpha, r \in \mathbb{R}^+$.
Output: Cipherstate $\rho \in \mathcal{D}(\mathcal{H}_1^{\otimes N})$.

1: Compute ciphertext $m' = \texttt{Enc}_{\text{base}}(s, m)$.
2: Compute codeword $c = \texttt{Encode}(m')$.
3: **for** $i = 1$ to N **do**
4: Prepare single mode squeezed state ρ_i with displacement d_i and covariance matrix Γ_i, where $d_i = [\alpha(-1)^{c_i} + k_i]\binom{\phi_i}{\phi_i}$, and $\Gamma_i = \begin{pmatrix} (\cosh r)^{2\phi_i - 1} & 0 \\ 0 & (\cosh r)^{1-2\phi_i} \end{pmatrix}$.
5: **end for**
6: $\rho = \rho_1 \otimes \cdots \otimes \rho_N$.
7: Output ρ.

Algorithm 3 Decryption (Dec)

Input: Key $(s, \phi, k) \in \{0,1\}^z \times \{0,1\}^N \times \mathbb{R}^{\mathbb{N}}$; cipherstate $\rho \in \mathcal{D}(\mathcal{H}_1^{\otimes N})$.
Output: Message $\hat{m} \in \{0,1\}^n$.

1: **for** $i = 1$ to N **do**
2: Perform homodyne measurement on ρ_i in the $\phi_i \cdot \frac{\pi}{2}$ direction, resulting in outcome $y_i \in \mathbb{R}$.
3: $\hat{c}_i = \frac{1}{2} - \frac{1}{2}\text{sign}(y_i - k_i)$.
4: **end for**
5: $\hat{c} = (\hat{c}_1, \dots, \hat{c}_N)$.
6: $\mu = \texttt{Decode}(\hat{c})$.
7: $\hat{m} = \texttt{Dec}_{\text{base}}(s, \mu)$.
8: Output $\hat{m}$.

Theorem 3.1. *Our construction is a QECM scheme with decryption failure parameter* $\varepsilon_{\text{DF}} = \exp\left[- N D_{\text{KL}}(\frac{t+1}{N} || \frac{1}{2}\text{Erfc}(\alpha\sqrt{\cosh r}))\right]$*, where* D_{KL} *stands for the binary Kullback-Leibler divergence,* $D_{\text{KL}}(a||b) = a \ln \frac{a}{b} + (1-a) \ln \frac{1-a}{1-b}$*, and Erfc is the complementary error function.*

Proof. It is trivial to see that the triplet (Key, Enc, Dec) fits the format of the CPTP maps in Def. 2.1. All that is left to show is that the protocol satisfies the correctness condition. A bit error in the codeword bit c_i occurs when $y_i - k_i$ has the wrong sign. Consider, without loss of generality, $c_i = 0$. Then the bit error probability is $\beta := \Pr[Y_i - k_i < 0]$. When k_i is known, the random variable $Y_i - k_i$ is gaussian-distributed with mean α and variance $1/(2\cosh r)$. Hence $\beta = \frac{1}{2} - \frac{1}{2}\text{Erf}(\alpha\sqrt{\cosh r})$.

A decoding error can occur only when there are more than t bitflips; this occurs with probability $P = \sum_{j=t+1}^{N} \binom{N}{j} \beta^j (1-\beta)^{N-j} = \sum_{\ell=0}^{N-t-1} \binom{N}{\ell} \beta^{N-\ell} (1-\beta)^{\ell}$. Using a Chernoff bound on the binomial tail, we obtain

$$P \leq \exp\Big[- N D_{\mathrm{KL}}(\frac{t+1}{N} || \beta) \Big]. \tag{11}$$

□

As an illustration, consider a ciphertext of size $N = 1000$. We set the protocol parameters $\alpha = 0.05$ and $r = 8$. This leads to a bit error probability of $\beta = 0.00317$. We choose $t = 16 (> N\beta = 3.16)$. This leads to an exponentially small failure probability $\varepsilon_{\mathrm{DF}} = 3.63 \times 10^{-7}$.

4 Proving Unclonability

4.1 Game hopping

We prove that our scheme satisfies the security definition of Def. 2.3. We do this as follows. (i) We rewrite the state preparation from the original prepare-and-send form to Entanglement Based (EB) form. (ii) We show, in a number of 'hops', that the cloning game that the security definition is based on is equivalent to a CV entanglement monogamy game. (iii) Finally we use Lemma 2.5 which upper bounds the winning probability in the CV entanglement monogamy game.

The sequence of games is as follows

- The cloning game for the actual QECM scheme.
- The cloning game for the Entanglement Based form of the QECM scheme.
- A variant of the above game, where now the keys s and k are *not* revealed to Bob and Charlie. Only the squeezing directions ϕ are revealed.
- The Full CV entanglement monogamy game.
- The Partial CV entanglement monogamy game (Game 2.4).

4.2 Entanglement based version

The first hop is to replace (in the encryption algorithm) the drawing of k_i and the preparation of the single-mode state ρ_i by the following procedure. (1) Prepare a two-mode entangled state $\rho^{(i)}_{ChA}$. (2) Do a homodyne measurement on the 'Ch' subsystem in the ϕ_i direction, yielding outcome u_i. Compute $k_i = [u_i - \alpha(-1)^{c_i}] \tanh r$.

The state $\rho^{(i)}_{ChA}$ must be such that it yields the correct distribution for k_i (truncated Gaussian) when the Ch system is measured in the ϕ_i direction, *and* the state of the 'A' subsystem gets projected to the correct squeezed coherent state. Without the truncation, $\rho^{(i)}_{ChA}$ would simply be given by the displaced EPR state mentioned in Section 2.2. In order to reproduce the truncated k_i-interval, however, the EPR state needs to be modified. The details are presented in Appendix A.

4.3 Unclonable Encryption game

The cloning attack (Def. 2.2) can be represented as the game below, between the Challenger and the three players $\mathcal{ABC}$.

Game A: Unclonable Encryption

1. **Initial state preparation.**
   The challenger Ch picks a random message m.
   Prepare-and-send: Ch runs Key (algorithm 1) and obtains (s, ϕ, k). He runs Enc (algorithm 2) on m. He sends the resulting cipherstate to Alice.
   Entanglement-based: Ch runs part of algorithm 1 and obtains (s, ϕ). He prepares an entangled state $\rho_{ChA} = \bigotimes_{i=1}^{N} \rho_{ChA}^{(i)}$, as explained in Section 4.2. He sends the 'A' part to Alice.

2. **Distributing quantum information to co-players.**
   Alice ($\mathcal{A}$), "splits" her quantum state and sends the two parts to her co-players, Bob ($\mathcal{B}$) and Charlie ($\mathcal{C}$). After that $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$ are no longer allowed to communicate.

3. **Key opening and response.**
   Prepare-and-send: Ch announces the key (s, ϕ, k). $\mathcal{B}$ and $\mathcal{C}$ respond with m_B and m_C respectively.
   Entanglement-based: Ch does homodyne measurements on his modes and computes k from the outcomes. Ch announces the key (s, ϕ, k). $\mathcal{B}$ and $\mathcal{C}$ respond with m_B and m_C respectively.

4. **Winning condition.**
   The triplet of players $\mathcal{ABC}$ win the game if $m_B = m_C = m$.

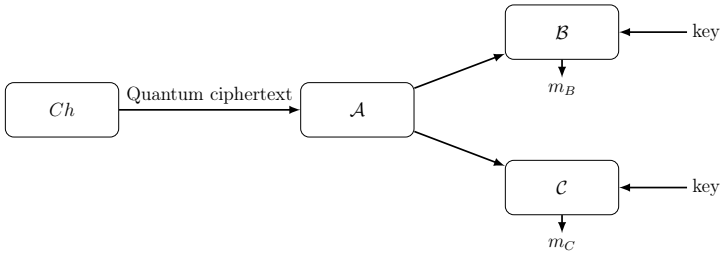

Figure 4.1: *Schematic representation of the unclonable encryption game.*

Note that in this game the measurement by Ch has been postponed to step 3, as opposed to step 1 in the direct EB description of the preparation of A's state.

Game B: EB Intermediate Unclonable Encryption

1. **Initial state preparation.**
   Ch prepares the entangled state ρ_{ChA}. He sends one half of the state to Alice.

2. **Distributing quantum information to co-players.**

3. **Key opening (only ϕ) and response**.
   Ch measures u and announces ϕ. $\mathcal{B}$ and $\mathcal{C}$ respond with u_B and u_C respectively. (Their guess for u.)

4. **Winning condition.**
   For $i \in \{1, \ldots, N\}$ Ch calculates $c_{Bi} = \frac{1}{2} - \frac{1}{2}\text{sign}(u_{Bi} - \frac{k_i}{\tanh r})$ and $c_{Ci} = \frac{1}{2} - \frac{1}{2}\text{sign}(u_{Ci} - \frac{k_i}{\tanh r})$. Then $m_B = \texttt{Dec}_{\text{base}}(s, \texttt{Decode}(c_B))$ and $m_C = \texttt{Dec}_{\text{base}}(s, \texttt{Decode}(c_C))$. The players $\mathcal{ABC}$ win if $m_B = m_C = m$.

Lemma 4.2. *Winning the Entanglement Based Unclonable Encryption game (Game A) is equivalent to winning the EB Intermediate Unclonable Encryption game (Game B) .*

Proof. Consider the EB version of Game A. Given k_i, Bob and Charlie know that there are only two possible values for u_i, namely $u_i = \frac{k_i}{\tanh r} \pm \alpha$. Hence, getting a bit c_i correct is equivalent to making the correct binary choice for u_i. From the fact that, by construction, $u_i \in (-2\alpha, 2\alpha)$, it follows that this in turn is equivalent to determining u_i with a resolution of α or better.
In Game B it is precisely such resolution on the u-axis that is required to guess c_i. □

4.4 Monogamy of entanglement game

We introduce a version of the entanglement monogamy game that is closer to the encryption scheme than Game 2.4.

Game C: CV Full Monogamy of Entanglement game

1. **Initial state preparation.**
   Bob and Charlie prepare an M-mode state ρ_{ABC} across three registers: one for Alice, one for Bob, and one for Charlie. After preparation, they are no longer allowed to communicate.

2. **Alice's measurement choice and outcomes.**
   Alice chooses quadrature directions $(\theta_i)_{i=1}^M$, $\theta_i \in \{0, \frac{\pi}{2}\}$, such that $|i : \ \theta_i = 0| = M/2$, i.e. each direction is chosen exactly $M/2$ times. She then does homodyne detection of mode i in direction θ_i, getting outcomes that we denote as

$$q_i \quad \text{if } \theta_i = 0, \qquad p_i \quad \text{if } \theta_i = \tfrac{\pi}{2}.$$

3. **Announcement and responses.**
   Alice announces $(\theta_i)_{i=1}^M$. Bob outputs a list $q_B \in \mathbb{R}^{M/2}$ containing his guesses for Alice's q-values and a list $p_B \in \mathbb{R}^{M/2}$ containing his guesses for Alice's p-values. Similarly, Charlie outputs q_C, p_C.

4. **Winning condition.**
   Bob and Charlie win the game if

$$\|\, q - q_B \|_\infty < \delta, \quad \|\, p - p_B \|_\infty < \delta, \quad \|\, q - q_C \|_\infty < \varepsilon, \quad \text{and} \quad \|\, p - p_C \|_\infty < \varepsilon. \tag{12}$$

Lemma 4.3. *Let $w^{\mathrm{UE}}(N,t)$ denote the winning probability for Game B. Let $w^{\mathrm{full}}(N-2t)$ be the winning probability for Game C with $\delta = \alpha$, $\varepsilon = \alpha$, $M = N - 2t$. It holds that*

$$w^{\mathrm{UE}}(N,t) \leq 2^{N-n}\, w^{\mathrm{full}}(N-2t). \tag{13}$$

Proof. In Game B, getting a bit c_i correct is equivalent to getting u_i correct within a distance α. This corresponds to setting $\delta = \alpha$ and $\varepsilon = \alpha$ in Game C.

Next, in Game B it is required to get at least $N - t$ bits of c correct. Let $e \in \{0,1\}^N$ denote an error pattern. Let w_e^{UE} be the probability of winning with precisely error pattern e. The bit errors can be arbitrarily distributed over the q part and the p part, which does not nicely fit the symmetric structure of the monogamy game. In order to obtain the symmetric structure we loosen our requirements a little, and allow a surplus of bit errors in one block (if any) to be balanced by additionally allowed bit errors in the other block. In the worst case, all bit errors are located in one block; this leads to an allowed $2t$ errors. We write

$$w^{\mathrm{UE}}(N,t) = \sum_{e\in\{0,1\}^N : |e| \leq t} w_e^{\mathrm{UE}} \leq \sum_{e\in\{0,1\}^N : |e| \leq t} w^{\mathrm{full}}(N-2t) = w^{\mathrm{full}}(N-2t) \left| e \in \{0,1\}^N : |e| \leq t \right|. \tag{14}$$

Finally we apply the Hamming bound for binary codes. □

4.5 Main theorem

Theorem 4.4. *Our QECM scheme is unclonable secure according to Def. 2.3, with*

$$\tau(\lambda) = \frac{N}{2} + \left(\frac{N}{2} - t\right)\log(1+2\alpha) + t + \frac{1}{2}\log e. \tag{15}$$

Proof. The cloning attack in Def. 2.3 has a success probability equal to the winning probability in Game A. By Lemma 4.2, this equals the winning probability of Game B. Next, by Lemma 4.3 this is upper bounded by $2^{N-n}\, w^{\text{full}}(N-2t)$. We use the fact that the full monogamy game C is harder (or equally hard) to win than the partial monogamy game 2.4. Thus we have $w^{\text{full}}(N-2t) \leq \sqrt{e}(\frac{1}{2}+\alpha)^{\frac{N}{2}-t}$. □

Remark: For $n \gg 1, \alpha \ll 1$ it holds that $\tau(\lambda) \approx \frac{N}{2}$. This is a large number; however, it still results in a winning probability that is exponentially small in n, which is the main objective for a QECM scheme.

Note that asymptotically N gets close to $\frac{n}{1-h(\beta)}$, where h is the binary entropy function and $\beta = \frac{1}{2}\text{Erfc}(\alpha\sqrt{\cosh r})$ is the bit error rate specified in the proof of Theorem 3.1.

We see from the result (15) that α needs to be smaller than approximately $\frac{1}{2}$, otherwise the attackers' winning probability is not exponentially small in n. This is unsurprising, for the following reason. One obvious cloning attack on our scheme is to perform a *heterodyne* measurement on each cipherstate mode individually. This corresponds to mixing the squeezed state with the vacuum state using a 50/50 beam splitter. The effect is a deterioration of the signal-to-noise ratio: On the one hand, the signal power α^2 is reduced to $\alpha^2/2$. On the other hand, the noise power goes from $\frac{1}{2\cosh r}$ to $\frac{1}{2}(\frac{1}{2} + \frac{1}{2\cosh r})$ due to the averaging with the vacuum's shot noise, which has power $\frac{1}{2}$. The resulting signal-to-noise ratio in the heterodyne measurement is $2\alpha^2/(1+\frac{1}{\cosh r}) \approx 2\alpha^2$. Hence at $\alpha = \mathcal{O}(1)$ there is significant leakage about the codeword bits c_i, putting the unclonability at risk.

It is interesting to note that the precise value of the squeezing parameter r has little influence on the unclonability property. The r has to be set such that the bit error rate β for the legitimate recipient is manageable, so that decryption succeeds with very high probability. I.e. the signal-to-noise ratio needs to be large enough. Other than that, there are no constraints on r.

Fig. 4.5 plots the winning probability of our UE game as a function of message size, including also the UE game for conjugate coding into qubit states [6]. The conjugate coding scheme first one-time pads a message m and then bitwise encodes the ciphertext into the standard basis or the Hadamard basis. Note that Fig. 4.5 is just an illustration for a single choice of parameters, and is not representative. Furthermore, a comparison of DV and CV is not as straightforward as directly comparing the winning probability in the cloning game.

5 Discussion

We have introduced the first Unclonable Encryption scheme with Continuous-Variable states, and have given a proof of unclonability in the game-based framework of [6]. This brings an interesting cryptographic primitive into the CV domain, making it potentially cheaper to implement.

We see several avenues for improvement. It is possible to get some extra tightness in the inequalities. The $\tau(\lambda)$ can be slightly improved by performing error correction on the q-block and p-block separately; this would change the $N-2t$ in (13) to $N-t$.

It would be interesting to see if a tighter upper bound can be derived for the winning probability in Game C (full monogamy game). We have now used the upper bound for the *partial* game, and we expect that this introduces a significant loss of tightness.

As a topic for future work we mention tolerance to channel noise. Channel noise will have a rather severe impact, since attenuation reduces the squeezing. In DV a similar effect arises: photon loss necessitates the use of error-correcting codes that can deal with erasures, which has a crippling effect on the efficiency of UE.

Furthermore, this work can pave the way for other schemes based on unclonability, for example, single-decryptor encryption [25], or revocable commitment.

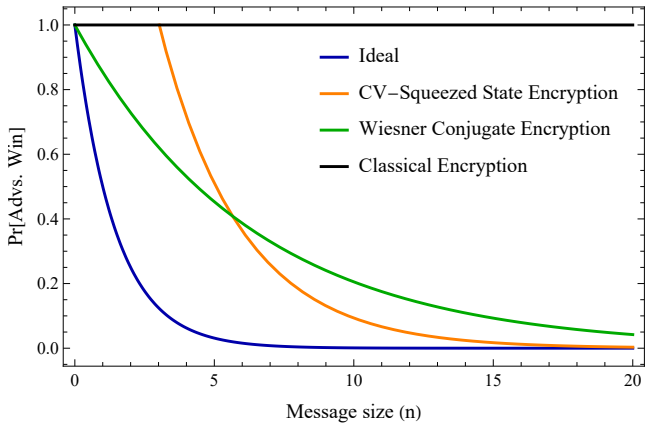

Figure 4.5: *Winning probabilities in the unclonability game for various schemes. The ideal curve is the simple guessing probability* $\left(\frac{1}{2}\right)^n$. *The wining probability of the conjugate coding scheme is at most* $\left(\frac{1}{2}+\frac{1}{2\sqrt{2}}\right)^n$. *For our CV scheme we plotted (15) with parameters* $r = 12, t = 1$ *and* $\alpha = 0.01$.

Appendix A The Entangled State

In the prepare-and-send scheme (algorithm 2), we restrict k_i to the interval $(-\alpha \tanh r, \alpha \tanh r)$. To ensure compatibility with the EB version, the corresponding outcome u_i of the homodyne measurement on the Challenger's mode needs to be similarly restricted, $u_i - \alpha(-1)^{c_i} \in (-\alpha, \alpha)$. Consider the Two-Mode Squeezed state without displacement

$$|\mathrm{EPR}_{r,\varphi}(0)\rangle = \frac{1}{\cosh\frac{r}{2}} \sum_{n=0}^{\infty} \left(-e^{i\varphi}\tanh\frac{r}{2}\right)^n |n,n\rangle. \tag{16}$$

In the position eigenbasis of both modes, the amplitude of this state can be represented using Hermite polynomials H_n. For $\phi = 0$ we have

$$\langle x_1, x_2|\mathrm{EPR}_{r,0}(0)\rangle = \frac{e^{-\frac{1}{2}(x_1^2+x_2^2)}}{\sqrt{\pi}\cosh\frac{r}{2}} \sum_{n=0}^{\infty} \frac{1}{n!} H_n(x_1)H_n(x_2)\left(\frac{1}{2}\tanh\frac{r}{2}\right)^n. \tag{17}$$

Using the identity $\sum_{n=0}^{\infty} \frac{1}{n!}H_n(x)H_n(y)\left(\frac{z}{2}\right)^n = \frac{1}{\sqrt{1-z^2}}\exp\left[\frac{2z}{1+z}xy - \frac{z^2}{1-z^2}(x-y)^2\right]$, we can simplify the expression to

$$\langle x_1, x_2|\mathrm{EPR}_{r,0}(0)\rangle = \frac{1}{\sqrt{\pi}} e^{-\frac{1}{4}e^r(x_1-x_2)^2} e^{-\frac{1}{4}e^{-r}(x_1+x_2)^2}. \tag{18}$$

Now we switch to the displaced EPR state, where both modes are translated over $\pm\alpha$ in the $\phi = 0$ direction,

$$\langle x_1, x_2|\mathrm{EPR}_{r,0}(\pm\alpha)\rangle = \frac{1}{\sqrt{\pi}} e^{-\frac{1}{4}e^r(x_1-x_2)^2} e^{-\frac{1}{4}e^{-r}(x_1+x_2\mp 2\alpha)^2}. \tag{19}$$

Finally we impose the restriction on the homodyne outcome by restricting the x_1-range,

$$|\mathrm{EPR}_{r,0}(\pm\alpha)\rangle_{\mathrm{restr}} \propto \int_{\pm\alpha-\alpha}^{\pm\alpha+\alpha} \mathrm{d}x_1\ |x_1\rangle \otimes \int_{-\infty}^{\infty} \mathrm{d}x_2\ |x_2\rangle\langle x_1, x_2|\mathrm{EPR}_{r,0}(\pm\alpha)\rangle \tag{20}$$

$$\propto \int_{\pm\alpha-\alpha}^{\pm\alpha+\alpha} \mathrm{d}x_1\ |x_1\rangle e^{-\frac{(x_1\mp\alpha)^2}{2\cosh r}} \otimes \int_{-\infty}^{\infty} \mathrm{d}x_2\ |x_2\rangle e^{-\frac{\cosh r}{2}[x_2\mp\alpha-(x_1\mp\alpha)\tanh r]^2} \tag{21}$$

$$= \int_{-\alpha}^{\alpha} \mathrm{d}x\ |x\pm\alpha\rangle e^{-\frac{x^2}{2\cosh r}} \otimes \int_{-\infty}^{\infty} \mathrm{d}x'\ |x'\pm\alpha\rangle e^{-\frac{\cosh r}{2}[x'-x\tanh r]^2}. \tag{22}$$

It is not necessary to have a Gaussian-like distribution for the quadrature in the Challenger's mode. Here we specified an entangled state that resembles the standard TMS state, because it has a special property: The effect of $|\mathrm{EPR}_{r,\phi}(\pm\alpha)\rangle_{\mathrm{restr}}$ can be obtained (although inefficiently) by taking the standard TMS state, homodyne measuring the Ch mode, and discarding if u_i does not lie in the right interval.

References

[1] D. Dieks. "Communication by EPR devices". In: *Physics Letters A* 92.6 (1982), pp. 271–272. ISSN: 0375-9601. DOI: https://doi.org/10.1016/0375-9601(82)90084-6. URL: https://www.sciencedirect.com/science/article/pii/0375960182900846.

[2] James L. Park. "The Concept of Transition in Quantum Mechanics". In: *Foundations of Physics* 1.1 (1970), pp. 23–33. DOI: 10.1007/BF00708652. URL: https://doi.org/10.1007/BF00708652.

[3] W. K. Wootters and W. H. Zurek. "A Single Quantum Cannot Be Cloned". In: *Nature* 299.5886 (1982), pp. 802–803. DOI: 10.1038/299802a0. URL: https://doi.org/10.1038/299802a0.

[4] Daniel Gottesman. "Uncloneable encryption". In: *Quantum Info. Comput.* 3.6 (Nov. 2003), pp. 581–602. ISSN: 1533-7146.

[5] Daan Leermakers and Boris Skoric. "Qubit-based unclonable encryption with key recycling". In: *Quant. Inf. Comput.* 21.11-12 (2021), pp. 901–930. DOI: 10.26421/QIC21.11-12-1.

[6] Anne Broadbent and Sébastien Lord. "Uncloneable Quantum Encryption via Oracles". In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Ed. by Steven T. Flammia. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. ISBN: 978-3-95977-146-7. DOI: 10.4230/LIPIcs.TQC.2020.4. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2020.4.

[7] Prabhanjan Ananth and Fatih Kaleoglu. "Unclonable Encryption, Revisited". In: *Theory of Cryptography*. Ed. by Kobbi Nissim and Brent Waters. Cham: Springer International Publishing, 2021, pp. 299–329. ISBN: 978-3-030-90459-3.

[8] Daan Leermakers and Boris Škorić. "Two-way unclonable encryption with a vulnerable sender". In: *International Journal of Quantum Information* 20.02 (2022), p. 2150037. DOI: 10.1142/S0219749921500374. eprint: https://doi.org/10.1142/S0219749921500374. URL: https://doi.org/10.1142/S0219749921500374.

[9] Srijita Kundu and Ernest Y.-Z. Tan. "Device-independent uncloneable encryption". In: *Quantum* 9 (2025). Preprint: arXiv:2210.01058v5, p. 1582. DOI: 10.22331/q-2025-01-08-1582. URL: https://quantum-journal.org/papers/q-2025-01-08-1582/.

[10] Prabhanjan Ananth et al. *On the Feasibility of Unclonable Encryption, and More*. Cryptology ePrint Archive, Paper 2022/884. 2022. URL: https://eprint.iacr.org/2022/884.

[11] Christian Majenz, Christian Schaffner, and Mehrdad Tahmasbi. "Limitations on Uncloneable Encryption and Simultaneous One-Way-to-Hiding". In: *IACR Cryptol. ePrint Arch.* 2021 (2021), p. 408. URL: https://api.semanticscholar.org/CorpusID:232379986.

[12] Timothy C. Ralph. "Continuous variable quantum cryptography". In: *Phys. Rev. A* 61 (1 Dec. 1999), p. 010303. DOI: 10.1103/PhysRevA.61.010303. URL: https://link.aps.org/doi/10.1103/PhysRevA.61.010303.

[13] Mark Hillery. "Quantum cryptography with squeezed states". In: *Phys. Rev. A* 61 (2 Jan. 2000), p. 022309. DOI: 10.1103/PhysRevA.61.022309. URL: https://link.aps.org/doi/10.1103/PhysRevA.61.022309.

[14] Margaret D. Reid. "Quantum cryptography with a predetermined key using continuous-variable Einstein-Podolsky-Rosen correlations". In: *Phys. Rev. A* 62 (6 Nov. 2000), p. 062308. DOI: 10.1103/PhysRevA.62.062308. URL: https://link.aps.org/doi/10.1103/PhysRevA.62.062308.

[15] Nicolas J. Cerf, Mel Lévy, and Gilles Van Assche. "Quantum distribution of Gaussian keys using squeezed states". In: *Phys. Rev. A* 63 (5 Apr. 2001), p. 052311. DOI: 10.1103/PhysRevA.63.052311. URL: https://link.aps.org/doi/10.1103/PhysRevA.63.052311.

[16] Frédéric Grosshans and Philippe Grangier. "Continuous Variable Quantum Cryptography Using Coherent States". In: *Phys. Rev. Lett.* 88 (5 Jan. 2002), p. 057902. DOI: 10.1103/PhysRevLett.88.057902. URL: https://link.aps.org/doi/10.1103/PhysRevLett.88.057902.

[17] Frédéric Grosshans et al. "Quantum key distribution using Gaussian-modulated coherent states". In: *Nature* 421 (Feb. 2003), pp. 238–41. DOI: 10.1038/nature01289.

[18] Frédéric Grosshans et al. "Virtual Entanglement and Reconciliation Protocols for Quantum Cryptography with Continuous Variables". In: *Quantum Info. Comput.* 3.7 (Oct. 2003), pp. 535–552. ISSN: 1533-7146.

[19] Christian Weedbrook et al. "Quantum Cryptography Without Switching". In: *Phys. Rev. Lett.* 93 (17 Oct. 2004), p. 170504. DOI: 10.1103/PhysRevLett.93.170504. URL: https://link.aps.org/doi/10.1103/PhysRevLett.93.170504.

[20] Raul Garcia-Patron Sanchez. "Quantum information with optical continuous variables: from Bell tests to key distribution". PhD thesis. Université libre de Bruxelles, 2007.

[21] Anthony Leverrier. "Theoretical study of continuous-variable quantum key distribution". Theses. Télécom ParisTech, Nov. 2009. URL: https://pastel.hal.science/tel-00451021.

[22] Eric Culf, Thomas Vidick, and Victor V. Albert. "Group coset monogamy games and an application to device-independent continuous-variable QKD". In: *ArXiv* abs/2212.03935 (2022). URL: https://api.semanticscholar.org/CorpusID:254408906.

[23] Christian Weedbrook et al. "Gaussian quantum information". In: *Rev. Mod. Phys.* 84 (2 May 2012), pp. 621–669. DOI: 10.1103/RevModPhys.84.621. URL: https://link.aps.org/doi/10.1103/RevModPhys.84.621.

[24] Stephen Wiesner. "Conjugate coding". In: *SIGACT News* 15.1 (Jan. 1983), pp. 78–88. ISSN: 0163-5700. DOI: 10.1145/1008908.1008920. URL: https://doi.org/10.1145/1008908.1008920.

[25] Marios Georgiou and Mark Zhandry. *Unclonable Decryption Keys*. Cryptology ePrint Archive, Paper 2020/877. 2020. URL: https://eprint.iacr.org/2020/877.

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

Deep Learning Based Lifetime Prediction on p-GaN Gate HEMTs

Zhixuan Ge[1], Shuoyan Zhao[1], Andrea Natale Tallarico[2], Maurizio Millesimo[2], Vladislav Volosov[2], Antonio Imbrugliasi[3], and Justin Dauwels[1]
[1]*Delft University of Technology, The Netherlands*
[2]*University of Bologna, Italy*
[3]*STMicroelectronics, Italy*

Abstract

Machinery health prognostics play an important role in maintenance by monitoring health conditions without frequent checks. Traditional health status checks may require machine downtime and, in some cases, could pose risks to certain components, potentially increasing operating costs. However, the health condition of machinery is often reflected in trends observed within monitored operational data. Estimating the Remaining Useful Life (RUL) from such data enables timely maintenance actions, thereby preventing unexpected failures and reducing associated costs and potential damage [1]. Although extensive research has been conducted on various areas of machinery, e.g., rolling bearings, turbofan engines, and lithium-ion batteries, the domain of power electronic devices, whose lifespans are affected by electrical and thermal operating conditions, has not yet received sufficient attention. This is particularly critical for power electronic systems used in applications such as power inverters and efficient energy conversion, where device reliability is paramount.

Therefore, this research attempts to fill this gap by implementing convolution networks with the Attention Mechanism for RUL prediction on the High Electron Mobility Transistor (HEMT) with a Schottky p-GAN gate, using data collected at the University of Bologna [2], under various environmental conditions. Specifically, we focus on p-GaN gate HEMTs, which are increasingly important for high-efficiency and high-frequency power applications. This work focuses on direct prediction of RUL, which is formulated as a regression problem with linearly decaying labels during operation. Owing to the environmental complexity, the lifetime of a GaN HEMT varies from a few seconds to several hours under accelerated stress conditions. The various combinations of temperature and stress voltage also resulted in diverse degradation processes, subsequently causing different degradation patterns. The pattern in the gate leakage current I_g serves as a critical indicator in assessing the degradation process ascribed in the gate region of the device, which is the most critical area under such a kind of reliability test, namely high temperature gate bias (HTGB). However, the prediction model can overestimate the importance of I_g by directly copying it as the predicted RUL or overfitting to fluctuations in I_g.

To overcome these unsolved complex problems, we implemented various time series processing architectures, combining Temporal Convolutional Network (TCN) with the Attention Mechanism to improve the long-term memory of the model, which enhances the prediction robustness against instability in measurements. The trend to copy I_g was mitigated by incorporating the device's elapsed operating time and extraction of temporal-domain information. As shown in Figure 1, the prediction rapidly settles close to the actual value and exhibits minor oscillations, which demonstrates a high prediction accuracy. In general, the network manages to predict RUL with high accuracy under non-extreme environmental conditions, guaranteeing reliable estimation of the power device's lifetime under typical operating conditions.

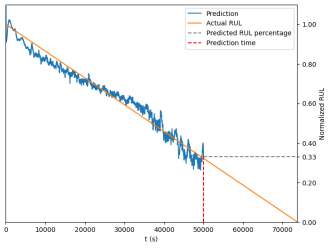

Figure 1. Process of dynamic RUL prediction at a certain time t = 50000 s. Prediction is close to actual RUL during operation.

Acknowledgments

This work, as a part of R-PODID project, is supported by the Chips Joint Undertaking and its members, including the top-up funding by National Authorities of Italy, Turkey, Portugal, The Netherlands, Czech Republic, Latvia, Greece, and Romania under grand agreement n° 101112338.

References

[1] Lei, Y., Li, N., Guo, L., Li, N., Yan, T., and Lin, J., "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," *Mechanical Systems and Signal Processing*, Vol. 104, 2018, p. 799–834.

[2] Tallarico, A., Posthuma, N., Bakeroot, B., Decoutere, S., Sangiorgi, E., and Fiegna, C., "Role of the AlGaN barrier on the long-term gate reliability of power HEMTs with p-GaN gate," *Microelectronics Reliability*, Vol. 114, 2020, p. 113872.

Differential Privacy for Markov Chains

Jasper Goseling
Mathematics of Operations Research
University of Twente, The Netherlands
j.goseling@utwente.nl

I. Introduction

Many aspects of the problem of publishing data while preserving privacy have been studied in the literature, but a largely open problem is how to sequentially publish data that is being collected over time. In this work we study this problem for the case that: i) data is generated according to a discrete-time Markov chain, ii) the goal is to publish the data itself (instead of e.g., aggregates), iii) at each time step, the data is published using the same privacy mechanism.

We present a bound on the overall privacy loss over the entire time horizon. This bound incorporates both the properties of the privacy mechanism itself and the privacy that is inherent due to Markov mixing. The main challenge is that the dependence between the data prevents the application of a basic strong data processing inequality.

II. Model

We consider a stationary, reversible and irreducible discrete-time Markov chain X_t, $t \in \mathbb{Z}$, on a finite state space $\mathcal{S}$. At each time step t we observe the state X_t and publish it using privacy mechanism $M : \mathcal{S} \to \mathcal{S}$. The privacy mechanism is assumed to be generalized randomized response, so that we have

$$D\left(M(x) \parallel M(x')\right) = \begin{cases} \epsilon, & \text{if } x \neq x', \\ 0, & \text{if } x = x', \end{cases} \tag{1}$$

where $D(\cdot\|\cdot)$ is the Kullback-Leibler divergence.

We denote the output as $Y_t = M(X_t)$ and $Y_{-n}^n = (Y_{-n}, Y_{-n+1}, \ldots, Y_n)$. Our goal is to bound the privacy loss on X_0. Therefore, let $P_{Y_{-n}^n|X_0=x}$ be the distribution of Y_{-n}^n given that $X_0 = x$. Our goal is bound the overall privacy loss, which we define as

$$\mathcal{L} = \lim_{n\to\infty} \max_{x,x'} D\left(P_{Y_{-n}^n|X_0=x} \,\middle\|\, P_{Y_{-n}^n|X_0=x'}\right). \tag{2}$$

We will express our results in terms of π, the stationary distribution of the Markov chain, and λ_2, the second largest eigenvalue of the transition matrix.

III. Results

Theorem 1. *The overall privacy loss $\mathcal{L}$ satisfies*

$$\mathcal{L} \leq \epsilon \left(1 + 2\sqrt{\max_x \frac{1}{\pi(x)}} \frac{e^{-(1-\lambda_2)}}{1 - e^{-(1-\lambda_2)}}\right). \tag{3}$$

An important part of the proof of this theorem is the following result, where P_U and P_V are arbitrary distributions and γ is a coupling of the two:

$$\begin{aligned} &D\left(M \circ P_U \| M \circ P_V\right) \\ &\quad = D\left(\sum_{u,v} \gamma(u,v) P_{M(u)} \middle\| \sum_{u,v} \gamma(u,v) P_{M(v)}\right) && (4) \\ &\quad \leq \sum_{u,v} \gamma(u,v) D\left(P_{M(u)} \middle\| P_{M(v)}\right) && (5) \\ &\quad \leq \epsilon \left\|P_U - P_V\right\|_{\mathrm{TV}}. && (6) \end{aligned}$$

In the above we use convexity of the KL-divergence after interpreting the distributions as mixtures with given components. Our coupling argument is similar to bounds appearing in [1].

IV. Discussion

Our use of the Kullback-Leibler divergence is somewhat uncommon in the context of differential privacy. The generalization to the α-Rényi divergence would be more appropriate [2], where the case of $\alpha \to \infty$ would give the most commonly used

$$\frac{\Pr(M(x) = y)}{\Pr(M(x') = y)} \leq e^{\epsilon}, \tag{7}$$

as the equivalent of (1). However, Rényi divergence does not satisfy the convexity property that we have used, so this is left as future work.

References

[1] F. Nielsen and R. Nock, "On w-mixtures: Finite convex combinations of prescribed component distributions," *arXiv preprint arXiv:1708.00568*, 2017.

[2] I. Mironov, "Rényi differential privacy," in *2017 IEEE 30th computer security foundations symposium (CSF)*, IEEE, 2017, pp. 263–275.

Reverse Engineering and FPGA Implementation of an OOK Receiver

M.A. Hemza, Z.J. Kohnen, A. Alvarado
Information and Communication Theory Lab
Signal Processing Systems Group
Department of Electrical Engineering, TU/e, The Netherlands
m.a.hemza@student.tue.nl, z.j.kohnen@student.tue.nl, a.alvarado@tue.nl

Abstract

On-Off Keying (OOK) modulation remains a widely adopted technology in low-cost wireless communication applications due to its simplicity and minimal hardware requirements. However, many wireless sensors, such as thermostats, rely on vendor-specific OOK-based signaling protocols, limiting integration into broader IoT frameworks.

This work presents two key contributions:

1. **Reverse Engineering Methodology for Proprietary Protocols:** We introduce an approach to decoding OOK-based wireless protocols without manufacturer specifications. The methodology consists of four steps: (i) capturing raw OOK waveform, (ii) analyzing timing to identify symbol patterns, (iii) inferring frame format and encoding and (iv) validating the decoded protocol.

2. **Digital FPGA-Based OOK Receiver:** We design and implement a noncoherent OOK receiver on an FPGA. The receiver comprises modules for envelope detection, frame synchronization, and data decoding, therefore eliminating the need for most of the analog front-end.

The experiment was set up as follows: A commercial wireless thermostat from Watts France transmitted OOK-modulated signals at 433MHz. The waveforms were captured using a PlutoSDR and downconverted to a lower intermediate frequency for processing. After protocol reverse engineering, a digital receiver was implemented on a Pynq Z2 development board, successfully decoding the downconverted waveforms. The design demonstrated reliable performance across multiple test scenarios, confirming both the correctness of the reverse engineered protocol and the functionality of the receiver.

By combining reverse engineering techniques with a receiver implementation, this work bridges the gap between closed sensor networks and open wireless systems. It facilitates greater interconnection in IoT applications and highlights the potential of digital OOK receivers, which still require initial analog downconversion, as alternatives to SDRs and microcontroller solutions.

Privacy Preserving Crowd Counting using Deep Learning on Range-Doppler Maps

Hippolyte Hilgers, Martin Willame, Gilles Monnoyer de Galland, Jérôme Louveaux, Christophe De Vleeschouwer, Anne-Sophie Collin
ICTEAM-ELEN UCLouvain
Louvain-la-Neuve, Belgium
hippolyte.hilgers@student.uclouvain.be

I. Motivation

Estimating crowd density in public space is essential for different applications such as safety monitoring during events like festivals or occupancy detection in smart buildings. Most state-of-the-art methods rely on camera-based systems combined with computer vision algorithms. While these approaches have proven to be effective, significant privacy concerns have been raised. To address this limitation, we propose a privacy-preserving alternative by performing crowd counting using a Frequency Modulated Continuous Wave (FMCW) radar. Specifically, we use Range-Doppler Maps (RDMs) derived from FMCW radar systems as input to a deep learning model, as traditional detectors like CFAR tend to fail in high-density scenarios, as demonstrated in [1].

II. Proposed Method

Building on prior work in radar-based crowd estimation [1], we designed a system that balances privacy and accuracy. We constructed a dataset of 50,000 RDMs, equivalent to approximately 50 minutes of recording, acquired via a FMCW radar placed alongside a camera in a large auditorium entrance hall. Each RDM is annotated with a pseudo Ground Truth (GT) label, obtained by applying a YOLOv3-based people detector [2] to synchronized camera frames. Camera data is used solely for labeling, not as model input. The annotated number of people per scene ranges from 0 to 25.

We propose a CNN-based method using a ResNet-18 architecture [3], trained from scratch on our RDM dataset. The network takes as input a stack of 10 consecutive RDMs captured over 625 ms and outputs a scalar estimate of the number of people in the scene.

III. Key Contributions

This work presents the following key contributions:

- The construction of a custom dataset composed of RDMs, annotated with pseudo ground truth labels derived from camera-based images.
- A deep learning framework based on a ResNet-18 architecture trained using a regression formulation to estimate crowd counting from RDM inputs.
- The use of 10 consecutive temporal RDMs as input to the model, which enhances prediction stability and mitigates the effects of signal noise.

IV. Results

As illustrated in Figure 1, the model estimates the number of people from RDMs, with a mean error of 1.9. Two main error sources were identified. First, under the supposition that static people are not detected on RDMs, the number of people is sometimes underestimated by the model. Second, pseudo labels from the YOLOv3 detector can be inaccurate, especially under occlusion. Despite these challenges, the approach shows strong potential for crowd counting using RDMs.

References

[1] L. Storrer, H. C. Yildirim, M. Willame, J. Louveaux, P. De Doncker, S. Pollin, and F. Horlin, "Crowd counting model training with the method of moments in electromagnetics," in *Proc. IEEE Joint Communications & Sensing (JC&S)*, 2023, pp. 1–6.
[2] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
[3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

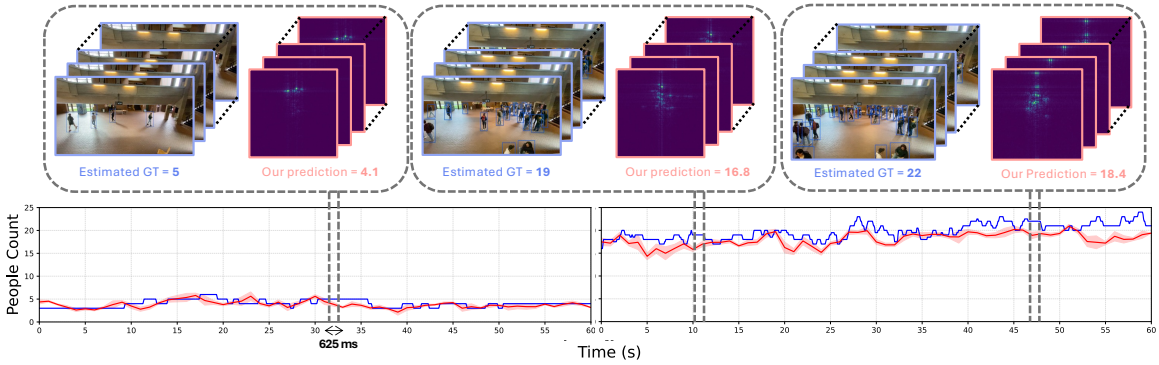

Fig. 1: Application of the method during two different 60-second acquisitions (GT vs. model prediction).

Numerical evaluation of Gaussian mixture entropy

Basheer Joudeh and Boris Škorić

We develop a method of approximation for the differential entropy $h(\mathbf{X})$ of a q-component Gaussian mixture random variable $\mathbf{X}$ in $\mathbb{R}^n$. We consider a random variable $\mathbf{X} \in \mathbb{R}^n$ whose probability density function is a Gaussian mixture with weights $\{p_i\}_{i=1}^q$. The Gaussian pdfs have covariance matrices $\{K_i\}_{i=1}^q$, and they are centered on points $\mathbf{w}_1, \ldots, \mathbf{w}_q \in \mathbb{R}^n$.

$$f_{\mathbf{X}}(\mathbf{x}) = \sum_{j=1}^{q} p_j \mathcal{N}_{\mathbf{w}_j, K_j}(\mathbf{x}) = (2\pi)^{-\frac{n}{2}} \sum_{j=1}^{q} p_j (\det K_j)^{-1/2} \exp\left\{ -\frac{1}{2}(\mathbf{x}-\mathbf{w}_j)^{\mathrm{T}} K_j^{-1} (\mathbf{x}-\mathbf{w}_j) \right\}. \quad (1)$$

Our method relies on finding coefficients $\{c_a\}_{a=1}^C$ such that:

$$-f_{\mathbf{X}} \ln \frac{f_{\mathbf{X}}}{m} \approx \sum_{a=1}^{C} c_a f_{\mathbf{X}}^a, \quad (2)$$

is a good approximation, and $m \propto \max_{\mathbb{R}^n} f_{\mathbf{X}}(\mathbf{x})$ can be chosen to enforce convergence. We find that the differential entropy $h(\mathbf{X})$ can be approximated by:

$$\bar{h}_{\vec{c},m}(\mathbf{X}) = -\ln m + \sum_{a=1}^{C} c_a \sum_{\substack{t_1,t_2,\ldots,t_q \geq 0 \\ t_1+t_2+\cdots+t_q=a}} \binom{a}{t_1, t_2, \ldots, t_q} \left(\prod_{i=1}^{q} p_i^{t_i} \right) D(\hat{t}), \quad (3)$$

where $D(\hat{t}) = \Pi_{j=1}^q \int \mathcal{N}_{\mathbf{w}_j,K_j}^{t_j}(\mathbf{x}) \mathrm{d}\mathbf{x}$. We implement equation (3) twice using two sets of coefficients that we obtain in different ways. The first is by use of the Taylor series of the logarithm, and the resulting approximation is denoted by $\bar{h}_{C,m}^{\text{Taylor}}(\mathbf{X})$. The second is by obtaining a polynomial fit approximation for the function $f(s) = -s \ln s$, and the resulting approximation for the differential entropy is denoted by $\bar{h}_{C,m}^{\text{Polyfit}}(\mathbf{X})$. The coefficients $\{c_a\}_{a=1}^C$ in both cases are given by:

$$c_a^{\text{Taylor}} = \begin{cases} H_{C-1}, & a = 1, \\ \dfrac{(-1)^{a+1}}{m^{a-1}} \dfrac{1}{a-1} \binom{C-1}{a-1}, & \text{otherwise.} \end{cases}, \qquad c_a^{\text{Polyfit}} = \frac{(M^{-1}\vec{z})_a}{m^{a-1}}, \quad (4)$$

where H_k is the k-th harmonic number, and matrix M and vector $\vec{z}$ have elements:

$$(M)_{ij} = \frac{2^{(i+j+r+1)}}{i+j+r+1}, \quad (\vec{z})_i = \frac{2^{(i+r+2)}(1-(i+r+2)\ln 2)}{(i+r+2)^2}. \quad (5)$$

We test both approximations numerically against the exact value of the differential entropy $h(\mathbf{X})$ for different Gaussian mixtures. We show that $\bar{h}_{C,m}^{\text{Taylor}}(\mathbf{X})$ provides an easy to compute lower bound to $h(\mathbf{X})$, while $\bar{h}_{C,m}^{\text{Polyfit}}(\mathbf{X})$ provides an accurate and efficient approximation to $h(\mathbf{X})$. $\bar{h}_{C,m}^{\text{Polyfit}}(\mathbf{X})$ is more accurate than known bounds, and conjectured to be much more resilient than perviously available approximations in the literature in high dimensions.

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

GNN-based Precoder Design and Fine-tuning for Cell-free Massive MIMO with Real-world CSI

Tianzheng Miao, Thomas Feys, Gilles Callebaut, Jarne Van Mulders, Emanuele Peschiera,
Md Arifur Rahman, François Rottenberg
KU Leuven, ESAT-WaveCore, Ghent Technology Campus, Ghent, Belgium
Research and Innovation Department IS-Wireless, Piaseczno, Poland

***Abstract*—Cell-free massive MIMO (CF-mMIMO) has emerged as a promising paradigm for delivering uniformly high-quality coverage in future wireless networks. To address the inherent challenges of precoding in such distributed systems, recent studies have explored the use of graph neural network (GNN)-based methods, using their powerful representation capabilities. However, these approaches have predominantly been trained and validated on synthetic datasets, leaving their generalizability to real-world propagation environments largely unverified. In this work, we initially pre-train the GNN using simulated channel state information (CSI) data, which incorporates standard propagation models and small-scale Rayleigh fading. Subsequently, we fine-tune the model on real-world CSI measurements collected from a physical testbed equipped with distributed access points (APs). To balance the retention of pre-trained features with adaptation to real-world conditions, we adopt a layer-freezing strategy during fine-tuning, wherein several GNN layers are frozen and only the later layers remain trainable. Numerical results demonstrate that the fine-tuned GNN significantly outperforms the pre-trained model, achieving an approximate 8.2 bits per channel use gain at $20\,\mathrm{dB}$ signal-to-noise ratio (SNR), corresponding to a $15.7\,\%$ improvement. These findings highlight the critical role of transfer learning and underscore the potential of GNN-based precoding techniques to effectively generalize from synthetic to real-world wireless environments.**

I. Introduction

Cell-free massive MIMO (CF-mMIMO) has emerged as a promising technology for future wireless communication systems, characterized by numerous distributed access points (APs) interconnected with one or more central-processing units (CPUs). This architecture enables cooperative service to all users within a given area via joint signal encoding and decoding based on users' channel state information (CSI), which is either locally estimated at each AP or centrally aggregated [1]. However, as the network scales up, efficiently managing the increasing volume of CSI and enabling robust precoding across diverse deployment scenarios remain key challenges. This highlights the need for scalable and generalizable precoding methods that can adapt to practical propagation conditions beyond idealized assumptions.

To overcome this challenge, recent advancements in learning-based methods, particularly those using graph neural networks (GNNs), have shown significant potential by exploiting the underlying topological structure of wireless networks [2]. For example, GNN-based methods have successfully facilitated implicit channel estimation in reflective intelligent surface (RIS) systems by directly mapping pilot signals to beamforming configurations through permutation-invariant architectures [3]. Additionally, GNN approaches have demonstrated adaptability against various sources of signal degradation. Recent studies indicate superior performance of GNN-based precoding compared to traditional methods, for instance in scenarios involving nonlinear power amplifier distortions [4].

However, despite the promising performance of existing GNN-based precoding approaches, most of them have been trained and evaluated on synthetic datasets. This is primarily due to the high cost and complexity involved in collecting large-scale, high-quality real-world channel measurements [5]. While synthetic data enables rapid experimentation, it often fails to capture the rich variability and hardware impairments present in practical deployments, limiting the generalization capability of learned models [6]. Ideally, training directly on real measurement data would improve robustness and deployment readiness. However, the limited availability of such data motivates the exploration of more efficient learning strategies. In this context, transfer learning (TL), which leverages knowledge gained from prior tasks or domains to improve learning efficiency in new ones [7], has emerged as a promising approach to bridge the sim-to-real gap. A typical example is to pre-train a model on synthetic data and then fine-tune it using a small amount of real-world measurements [8]. This motivates our work to investigate how GNN-based precoding models can be effectively adapted to realistic propagation scenarios using limited real measurements.

This work proposes a novel GNN-based precoding framework designed for CF-mMIMO systems. In the first stage, a GNN is designed to learn a mapping from channel matrices to precoding matrices in an unsupervised manner. The model is initially pretrained on a large-size synthetic dataset, where channel matrices are generated using standard geometric propagation models combined with small-scale Rayleigh fading. In the second stage, the pretrained GNN is fine-tuned using real-world CSI measurements collected from the Techtile testbed, enabling the model to adapt to practical channel conditions. To balance the retention of useful knowledge acquired during synthetic pretraining with the need to adapt to domain-specific distributions of real-world CSI, a strategic

This work was supported by the European Union's Horizon 2022 research and innovation program under Grant Agreement No 101120332. *(Corresponding author: Tianzheng Miao)*

layer-freezing scheme is introduced. In this scheme, selected layers of the network are frozen during fine-tuning to preserve generalizable representations while allowing deeper layers to specialize to the target domain. The performance of the fine-tuned model is evaluated on both synthetic and real CSI datasets, and benchmarked against conventional precoding techniques. Numerical results show that fine-tuning leads to a substantial performance gain—improving the sum-rate by approximately $8.2\,\mathrm{bits/channel, use}$, or about $15.7\,\%$—thus demonstrating the effectiveness of transfer learning in enabling practical generalization for real-world deployment.

Notations: Boldface lowercase and uppercase letters denote vectors and matrices, respectively. The operators $(\cdot)^{\mathrm{T}}$ and $(\cdot)^{H}$ indicate matrix transpose and conjugate transpose operations, respectively. The trace of a matrix is given by $\mathrm{Tr}(\cdot)$. The set of complex numbers is represented by $\mathbb{C}$.

II. System Model

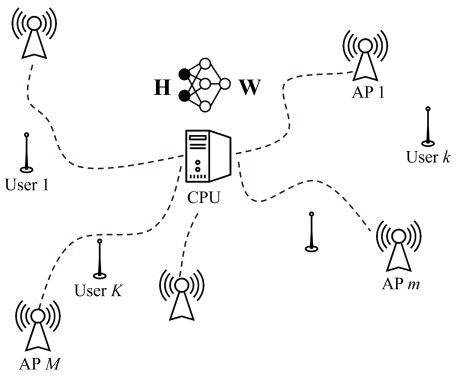

Fig. 1: System model of the considered CF-mMIMO network. Distributed APs serve multiple single-antenna user equipments (UEs) with coordination by a CPU, which handles joint signal processing.

As illustrated in Fig. 1, we consider a CF-mMIMO network comprising M single-antenna APs and K single-antenna UEs. Let $m = 1, 2, \ldots, M$ and $k = 1, 2, \ldots, K$ index the APs and UEs, respectively. All APs are equipped with a single isotropic antenna and connected via fronthaul links to a CPU, which performs centralized channel estimation and precoding based on global CSI.

In this work, we consider a fully centralized downlink scenario in which the APs are randomly located within a specific geographical area, and each UE is simultaneously served by all APs with $M > K$. The channel between the APs and the UEs is represented by the channel matrix $\mathbf{G} \in \mathbb{C}^{K \times M}$, where the channel coefficient between AP m and UE k is expressed as

$$g_{m,k} = \sqrt{\beta_{m,k}} h_{m,k} \tag{1}$$

where $\beta_{m,k}$ is the large-scale fading coefficient, capturing path loss effects, and $h_{m,k}$ is the small-scale fading coefficient. Specifically, the path loss (in dB) is modeled according to the Indoor Hotspot (InH) Non-Line-of-Sight (NLOS) scenario [9]

$$\beta_{m,k} = 32.4 + 31.9 \log_{10}(d_{m,k}) + 20 \log_{10}(f_c) \tag{2}$$

where d_{mk} denotes the distance in meters between AP m and UE k, and f_c is the carrier frequency in GHz. The small scale fading coefficients are independently and identically distributed (i.i.d.), where $h_{m,k} \sim \mathcal{CN}(0, 1)$ meaning that each coefficient follows a complex Gaussian distribution. To simultaneously serve all UEs, each AP participates in linear precoding. Let $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_K] \in \mathbb{C}^{M \times K}$ denote the precoding matrix, where $\mathbf{w}_k \in \mathbb{C}^M$ represents the precoding vector towards UE k from all APs. Accordingly, the received signal at UE k is thus given by

$$y_k = \mathbf{g}_k^{\mathrm{T}} \mathbf{w}_k s_k + \sum_{l=1, l \neq k}^{K} \mathbf{g}_k^{\mathrm{T}} \mathbf{w}_l s_l + n_k \tag{3}$$

where $\mathbf{g}_k \in \mathbb{C}^M$ denotes the channel vector between UE k and the APs, and $n_k \sim \mathcal{CN}(0, \sigma^2)$ is additive white Gaussian noise at UE k. The transmitted symbols $s_k \sim \mathcal{CN}(0, 1)$ are independent and identically distributed (i.i.d.) complex Gaussian random variables, uncorrelated across different UEs. Accordingly, the signal-to-interference-plus-noise ratio (SINR) at UE k is calculated as

$$\mathrm{SINR}_k = \frac{\left|\mathbf{g}_k^{\mathrm{T}} \mathbf{w}_k\right|^2}{\sum_{l=1, l \neq k}^{K} \left|\mathbf{g}_k^{\mathrm{T}} \mathbf{w}_l\right|^2 + \sigma^2}.$$

Thus, the sum rate of the system, combining the individual user rates, can be expressed as

$$R_{\text{sum}} = \sum_{k=1}^{K} R_k = \sum_{k=1}^{K} \log_2 \left(1 + \mathrm{SINR}_k\right). \tag{4}$$

III. GNN-based Precoder Design

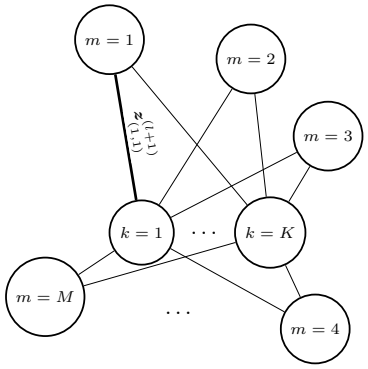

Fig. 2: Illustration of the graph representing the cell-free system

In the following, we describe the design and fine-tuning strategy of the proposed GNN-based precoder, shown in Fig. 2. The objective of our neural network is to learn a mapping from the channel matrix $\mathbf{G}$ to the corresponding precoding

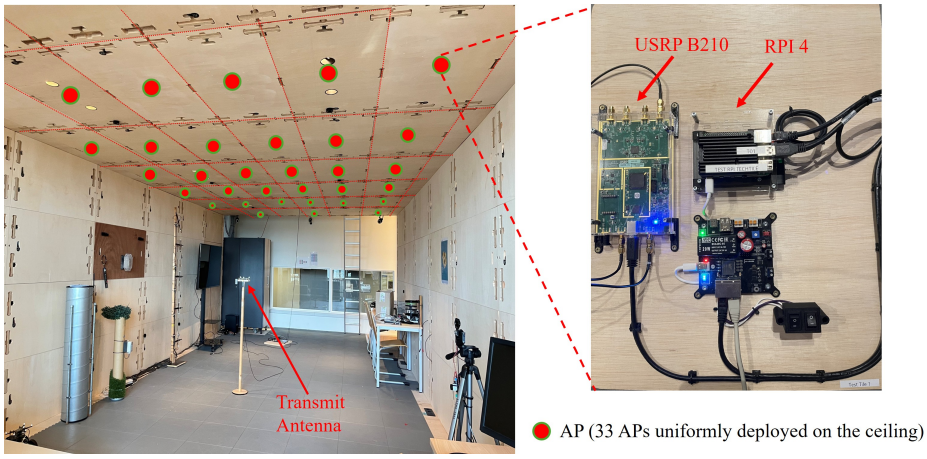

Fig. 3: Testbed setup (*Left:* Techtile environment with 33 APs on the ceiling and UE at floor. *Right:* Illustration of the hardware setup for each AP, deployed on the backside of the ceiling planks.)

matrix $\mathbf{W}$. Our GNN architecture comprises 8 layers, each executing message-passing operations on the graph-structured representation of the wireless network.

Inspired by recent advances in the literature [4], we adopt an edge-centric representation to better capture the wireless propagation characteristics, as edges naturally correspond to the communication channels between antennas at the APs and the UEs. In this formulation, the CF-mMIMO network is modelled as a bipartite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the vertex set $\mathcal{V}$ includes nodes representing the APs and the UEs. The edge set $\mathcal{E}$ encodes the wireless links, with CSI serving as edge attributes.

Each GNN layer updates its edge representations through message passing, aggregating information from neighboring nodes and edges according to

$$\mathbf{z}_{(m,k)}^{(l+1)} = \text{UPDATE}\left(\mathbf{z}_{(m,k)}^{(l)}, \mathbf{m}_{(m)}^{(l)}, \mathbf{m}_{(k)}^{(l)}\right) \tag{5}$$

where $\mathbf{z}_{(m,k)}^{(l)}$ denotes the representation of the edge connecting AP m and UE k at layer l, and $\mathbf{m}_{(m)}^{(l)}$, $\mathbf{m}_{(k)}^{(l)}$ represent aggregated messages from neighboring edges. Node messages are computed using aggregation functions as follows

$$\mathbf{m}_{(v)}^{(l)} = \text{AGGREGATE}\left(\mathbf{z}_{(v,u)}^{(l)} \mid u \in \mathcal{N}(v)\right). \tag{6}$$

In our model, the AGGREGATE function is defined as the element-wise mean over incoming edge features

$$\mathbf{m}(v)^{(l)} = \frac{1}{|\mathcal{N}(v)|} \sum_{u \in \mathcal{N}(v)} \mathbf{z}_{(v,u)}^{(l)}. \tag{7}$$

The UPDATE function linearly combines the current edge embedding with the aggregated messages from both endpoint nodes

$$\mathbf{z}_{(m,k)}^{(l+1)} = \sigma\left(\mathbf{W}_{\text{edge}}^{(l)} \mathbf{z}_{(m,k)}^{(l)} + \mathbf{W}_{m}^{(l)} \mathbf{m}_{(m)}^{(l)} + \mathbf{W}_{k}^{(l)} \mathbf{m}_{(k)}^{(l)}\right), \tag{8}$$

where $\sigma(\cdot)$ denotes a LeakyReLU activation function, and $\mathbf{W}_{\text{edge}}^{(l)}$, $\mathbf{W}_{m}^{(l)}$, and $\mathbf{W}_{k}^{(l)}$ are trainable weight matrices at layer l.

To adhere to transmit power constraints inherent in wireless communication systems, a power normalization step is integrated, defined by

$$\mathbf{W}^{\text{norm}} = \alpha \mathbf{W}, \quad \text{with} \quad \alpha = \sqrt{P_{\text{T}} / \text{Tr}(\mathbf{W}\mathbf{W}^{H})} \tag{9}$$

where P_{T} represents the total available transmit power. The proposed GNN-based precoder is trained in an unsupervised manner with the objective of maximizing the sum rate in Eq. (4).

IV. Transfer Learning with Real-World Data

In this study, we employ TL to transfer knowledge learned from synthetic data to real-world scenarios. Specifically, after an initial pretraining phase on large-scale synthetic datasets, we fine-tune the pretrained GNN model using a limited-size real-world dataset. This fine-tuning process is also conducted in an unsupervised manner.

A. Data Collection and Dataset Preparation

To assess the model's ability to generalize from simulated to real-world environments, we collected real CSI data using the Techtile testbed [10], which emulates a physical cell-free massive MIMO system. As illustrated in Fig. 3, our experimental setup includes 33 ceiling-mounted APs, each implemented using a universal software radio peripheral (USRP) B210 software-defined radio and managed via a dedicated Raspberry Pi (RPi) 4. Within the same space, a single UE was used to collect channel data.

During the data acquisition process, the UE transmits uplink pilot signals, which are coherently received by all APs to form the composite CSI. To ensure diverse spatial sampling and comprehensive coverage, the UE was moved to a different

position after each measurement. This procedure results in a dataset denoted as

$$\mathcal{H} = \{\mathbf{h}_1, \mathbf{h}_2, \ldots, \mathbf{h}_{500}\}$$

where each measurement vector $\mathbf{h}_i \in \mathbb{C}^{1\times M}$ represents the uplink CSI from a single-antenna user located at a specific spatial position to all $M = 33$ APs. This ensures spatial variability across 500 unique positions.

To simulate a two-user communication scenario, we construct a new dataset by pairing the single-user channel vectors. For each unordered pair of distinct indices (i, j) where $1 \leq i < j \leq 500$, we define a two-user sample as the tuple $(\mathbf{h}_i, \mathbf{h}_j)$. The total number of such combinations is given by the binomial coefficient $\binom{500}{2}$, resulting in 124 750 unique two-user channel instances. Each sample represents a realistic communication scenario in which two spatially separated users are jointly served by the same set of access points.

It is also possible to extend this setup to a four-user scenario by generating all unique unordered 4-tuples from the 500 single-user samples, with the total number of such combinations given by the binomial coefficient $\binom{500}{4}$. As this results in over 2.5 billion combinations, which is computationally infeasible to process in full. To balance the need for maintaining high-quality channel conditions with the requirement of controlling the dataset size for computational tractability and fair comparison across scenarios, we select the top 44 single-user samples based on channel strength $\|\mathbf{h}_i\|$. This selection ensures a sufficiently large number of unique four-user combinations, as $\binom{44}{4} = 135751 \geq 124750$. From these, we randomly sample 124 750 combinations without replacement to match the size of the two-user dataset.

For model training and evaluation, both the two-user and four-user datasets are partitioned into training, validation, and testing subsets, containing 80%, 10%, and 10% of the total samples, respectively. This split ensures sufficient diversity for effective model learning while maintaining reliable evaluation performance across unseen data. The complete preprocessed dataset is publicly available at Real-world CSI Dataset.

B. Fine-Tuning Strategy

A key challenge in deploying models trained on synthetic data is the distribution mismatch between simulated and real-world environments, commonly known as covariate shift [11]. In our context, real CSI exhibits non-idealities and measurement noise absent from simulated data, often leading to degraded performance when directly applying a pretrained model.

To effectively adapt the pretrained model to real-world data while retaining its learned representations, we adopt a layer-freezing approach during fine-tuning. The underlying GNN architecture consists of 8 layers, which allows for selective freezing at different depths of the network. Specifically, the first *l* layers of the model are kept fixed (i.e., their parameters are not updated), and the remaining *8-l* layers are retrained using the real-world dataset.

The proposed method aims to keep a balance between preserving useful knowledge acquired from simulation and enabling flexible adaptation to the domain shift introduced by real CSI. Furthermore, freezing early layers reduces the number of trainable parameters, which is particularly beneficial when only limited real training samples are available, as it helps mitigate overfitting while still allowing sufficient capacity in later layers to avoid underfitting.

To identify the optimal freezing point, we conducted an exhaustive evaluation across all possible values of $l \in \{0, 1, \ldots, 8\}$. Notably, Freeze_*0* corresponds to full fine-tuning, where all layers are updated, while Freeze_*8* denotes a fully frozen model, effectively identical to the pretrained network without any adaptation. This layer-wise investigation enables us to assess the trade-off between stability and adaptability under real deployment conditions.

V. Results

This section presents the results of our experiments, focusing on identifying the optimal layer-freezing strategy and evaluating the performance of the proposed GNN-precoder and fine-tuning network across various datasets. The signal-to-noise ratio (SNR) values shown on the x-axis are defined as $\mathrm{SNR}_{\mathrm{Tx}} = 10\log_{10}(P_t/\sigma^2)$, where σ^2 denotes the noise variance. This represents the transmit-side SNR prior to any channel fading or path loss effects.

A. Freezing Strategy and Fine-Tuning Performance

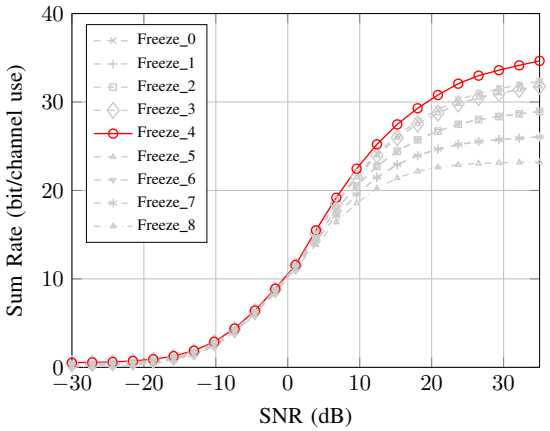

Fig. 4: Comparison of different freezing strategies. Freeze_*l* indicates that the first *l* layers of the network are frozen, while the remaining layers are retrained using the collected real-world dataset. Note that the evaluation was done on the real-world CSI dataset

The pretrained GNN was initially trained using a synthetic dataset containing 500 000 training samples and 50 000 validation samples, followed by testing on 10 000 samples. The training employed a learning rate of 0.005 over 20 epochs. Identical dataset sizes were used for both the two-user and four-user scenarios to ensure consistency in evaluation.

To determine the most effective freezing strategy for subsequent fine-tuning, we evaluated various configurations using real-world CSI data from a four-user scenario as the test set, with the learning rate kept consistent with that used during pretraining. As illustrated in Fig. 4, freezing the first 4 layers achieves the best performance. This configuration strikes a balance between retaining the pretrained model's prior knowledge and maintaining sufficient adaptability to the new real-world dataset. Notably, the performance of Freeze_*0* ranks second, slightly outperforming Freeze_*1* and Freeze_*7*. In contrast, Freeze_*8* yields the poorest performance, as freezing all layers prevents the model from adapting to the new data. Based on these findings, we adopted the Freeze_*4* strategy for all subsequent fine-tuning experiments.

B. Generalization on Real vs. Synthetic CSI

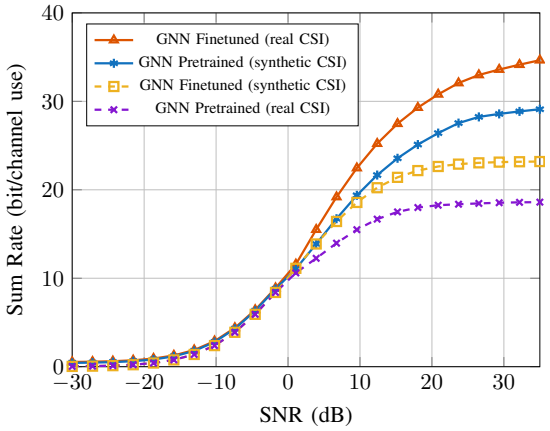

Fig. 5: Sum-rate performance comparison for different precoding methods with 4 UEs, evaluated on real and synthetic CSI.

Figure 5 illustrates the performance of the proposed GNN-based precoding network under various conditions. Both real and synthetic CSI datasets were used to evaluate the pretrained and fine-tuned models across a range of transmit SNR values. The results show that the fine-tuned model consistently outperforms the pretrained one on the real-world dataset, while it underperforms on the synthetic dataset compared to the pretrained model. This outcome reflects the effect of TL that fine-tuning enhances the model's ability to generalize to real data at the cost of some performance degradation on the synthetic domain. Moreover, the fine-tuned model performs better when evaluated on real CSI data than on synthetic data, demonstrating the benefits of domain adaptation through fine-tuning. Conversely, the pretrained model, which has not been exposed to real-world data, exhibits degraded performance when tested on the real CSI dataset, further highlighting the necessity of fine-tuning.

C. Scalability of Fine-Tuned GNN vs. Different Precoding Methods

This subsection evaluates the scalability of the proposed fine-tuned GNN model by comparing its performance to traditional precoding schemes in both two-user and four-user settings. As baselines, we adopt zero-forcing (ZF) and conjugate beamforming (CB), two classical and widely used precoding methods [12]. These methods are defined as follows

$$\mathbf{W} = \begin{cases} \alpha \mathbf{H}^H & \text{CB} \\ \alpha \mathbf{H}^H \left(\mathbf{H}\mathbf{H}^H\right)^{-1} & \text{ZF} \end{cases} \tag{10}$$

where ZF precoding is used as a performance reference in our experiment due to its theoretical ability to eliminate inter-user interference under ideal conditions.

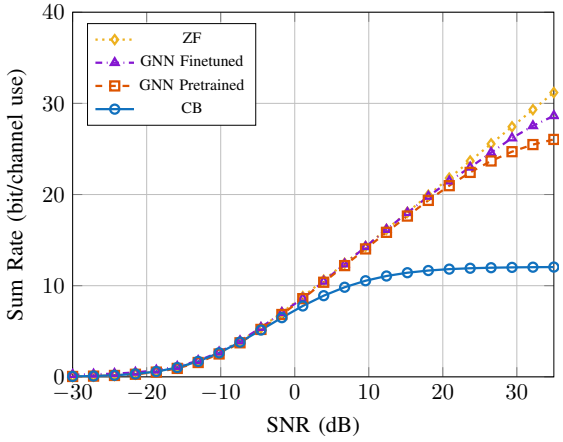

(a) Sum rate vs. SNR with 2 UEs

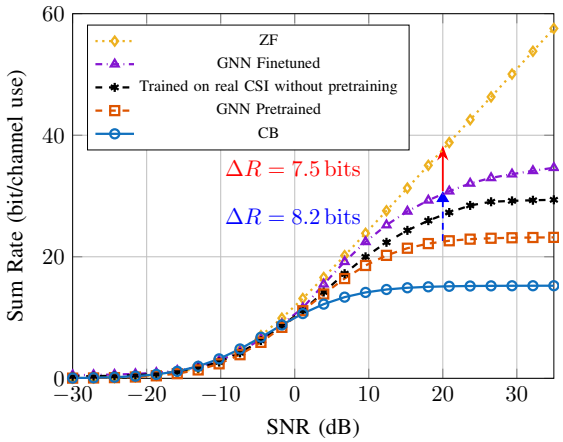

(b) Sum rate vs. SNR with 4 UEs

Fig. 6: Sum-rate performance of the CF-mMIMO system with different numbers of UEs ($M = 33$) and precoding schemes. All methods are evaluated on real-world CSI data.

As shown in Fig. 6(a), the pretrained GNN surpasses CB and approaches the performance of ZF up to $20\,\mathrm{dB}$ SNR. Beyond this point, its performance saturates. This limitation is attributed to the fact that the model was trained at a fixed SNR level, which reduces its ability to generalize to higher SNR regimes not seen during training. In contrast, ZF maintains robust performance by analytically computing the pseudo-inverse of the channel matrix.

In the four-user case (Fig. 6(b)), the pretrained GNN still outperforms CB, but lags significantly behind ZF, with a performance gap of $15.7\,\mathrm{bits/channel\,use}$. After applying fine-tuning, this gap is reduced to $7.5\,\mathrm{bits/channel\,use}$, corresponding to a relative improvement of approximately $15.7\,\%$. The increased number of users introduces more complex inter-user interference, which challenges the generalization ability of the pretrained model. However, the fine-tuning process enables the GNN to adapt to these more difficult scenarios by learning from real-world interference patterns, thereby significantly narrowing the performance gap.

The performance of the fine-tuned GNN compared with the baseline model trained directly on real CSI without pretraining is shown in Fig. 6(b). It can be observed that the baseline model fails to match the performance of the fine-tuned GNN. This performance gap is primarily due to the baseline being initialized with random weights, which can lead to convergence toward suboptimal solutions. In contrast, fine-tuning benefits from a favorable initialization derived from pretraining, which guides the optimization toward better local minima. Additionally, the inferior performance of the baseline can be attributed to the limited size of the real dataset, which may be insufficient to capture the full distributional characteristics of the underlying wireless channel.

VI. Conclusion

In this paper, we proposed a graph neural network (GNN)-based precoding framework for cell-free massive MIMO (CF-mMIMO) systems, with a focus on enhancing its practical deployment through the application of transfer learning (TL). The model was initially pretrained on a large-size synthetic channel state information (CSI) dataset generated using standard geometric propagation models combined with small-scale Rayleigh fading. To enable real-world applicability, the pretrained GNN was subsequently fine-tuned using measured CSI data collected from a physical testbed featuring distributed access points (APs). To balance knowledge retention from pretraining and adaptability to domain shifts in real-world environments, a strategic layer-freezing scheme was employed during fine-tuning. Experimental results demonstrate that the fine-tuned GNN achieves a notable performance improvement over the pretrained model, increasing the sum-rate by approximately $8.2\,\mathrm{bits/channel\,use}$, or about $15.7\,\%$.

References

[1] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free massive MIMO versus small cells," *IEEE Transactions on Wireless Communications*, vol. 16, no. 3, pp. 1834–1850, 2017.

[2] M. Lee, G. Yu, H. Dai, and G. Y. Li, "Graph neural networks meet wireless communications: Motivation, applications, and future directions," *IEEE Wireless Communications*, vol. 29, no. 5, pp. 12–19, 2022.

[3] T. Jiang, H. V. Cheng, and W. Yu, "Learning to Reflect and to Beamform for Intelligent Reflecting Surface With Implicit Channel Estimation," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 7, pp. 1931–1945, Jul. 2021.

[4] T. Feys, L. Van der Perre, and F. Rottenberg, "Toward energy-efficient massive MIMO: Graph neural network precoding for mitigating non-linear PA distortion," *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 1, pp. 184–201, 2025.

[5] A. Alkhateeb, "DeepMIMO: A Generic Deep Learning Dataset for Millimeter Wave and Massive MIMO Applications," Feb. 2019.

[6] Y. Huangfu, J. Wang, S. Dai, R. Li, J. Wang, C. Huang, and Z. Zhang, "WAIR-D: Wireless AI research dataset," *arXiv preprint arXiv:2212.02159*, 2022.

[7] M. Wang, Y. Lin, Q. Tian, and G. Si, "Transfer Learning Promotes 6G Wireless Communications: Recent Advances and Future Challenges," *IEEE Transactions on Reliability*, vol. 70, no. 2, pp. 790–807, Jun. 2021.

[8] ——, "Transfer Learning Promotes 6G Wireless Communications: Recent Advances and Future Challenges," *IEEE Transactions on Reliability*, vol. 70, no. 2, pp. 790–807, Jun. 2021.

[9] 3GPP, "Study on channel model for frequencies from 0.5 to 100 ghz (release 17)," 3GPP, Tech. Rep. TR 38.901 V17.0.0, 2022, (March 2022).

[10] G. Callebaut, J. Van Mulders, G. Ottoy, D. Delabie, B. Cox, N. Stevens, and L. Van der Perre, "Techtile: Open 6G R&D testbed for communication, positioning, sensing, WPT and federated learning," in *Proc. Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 2022, pp. 417–422.

[11] J. Quinonero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, *Dataset Shift in Machine Learning*. MIT Press, 2008.

[12] E. Björnson, L. Sanguinetti, J. Hoydis, and M. Debbah, "Optimal Design of Energy-Efficient Multi-User MIMO Systems: Is Massive MIMO the Answer?" *IEEE Transactions on Wireless Communications*, vol. 14, no. 6, pp. 3059–3075, Jun. 2015.

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

Joint Pilot and Data-based Localization Exploiting Direct and Relative Information Between Receivers

Mathieu Reniers*, Martin Willame*, Laurence Defraigne*, Gilles Monnoyer*, Jérôme Louveaux*, Luc Vandendorpe*,
*UCLouvain - Université Catholique de Louvain, Louvain-La-Neuve, Belgium.

I. Context and Contributions

The new generations of Wi-Fi and mobile communication systems aim to support multiple services, including joint communication and sensing (JCAS). However, these functions present conflicting objectives: communication seeks to reduce the pilot-to-data ratio for increased data rates while maintaining sufficient pilots for channel estimation, whereas positioning aims to maximize this ratio to improve localization performance. We explore how previously neglected [1] unknown data can be combined with pilots to enhance localization in JCAS scenarios. Our **contributions** can be summarized as follows. We propose two novel approaches leveraging jointly the pilots and data symbols: a **projection**-based estimator considering data estimation without any constellation information, and **transformation**-based estimators incorporating or not a priori constellation information on data symbols. We demonstrate the benefits of both methods over traditional approaches through numerical simulations.

II. System Model

A single user located in $\boldsymbol{x}$ transmits an uplink OFDM signal with P pilot symbols and D data symbols to N single-antenna receiver nodes, located in $\{\boldsymbol{x}_n\}_{n=1}^{N}$. Assuming perfect time and phase synchronization, the received pilots and data signals on a single subcarrier are expressed as

$$\boldsymbol{Y}_{\mathrm{P}} = \boldsymbol{a}(\boldsymbol{x})\boldsymbol{s}_{\mathrm{P}}^{T} + \boldsymbol{N}_{\mathrm{P}} \quad \in \mathbb{C}^{N\times P}, \tag{1}$$

$$\boldsymbol{Y}_{\mathrm{D}} = \boldsymbol{a}(\boldsymbol{x})\boldsymbol{s}_{\mathrm{D}}^{T} + \boldsymbol{N}_{\mathrm{D}} \quad \in \mathbb{C}^{N\times D}, \tag{2}$$

with known BPSK pilots symbols $\boldsymbol{s}_{\mathrm{P}} \in \mathcal{C}_{\mathrm{P}}^{P} \subset \mathbb{C}^{P}$, unknown M-QAM data symbols $\boldsymbol{s}_{\mathrm{D}} \in \mathcal{C}_{\mathrm{D}}^{D} \subset \mathbb{C}^{D}$, and steering vector

$$\boldsymbol{a}(\boldsymbol{x}) = \left[\exp\left(-j2\pi\tfrac{d(\boldsymbol{x},\boldsymbol{x}_1)}{\lambda}\right)\cdots\exp\left(-j2\pi\tfrac{d(\boldsymbol{x},\boldsymbol{x}_N)}{\lambda}\right)\right]^{T} \tag{3}$$

$\in \mathbb{C}^{N}$ with $d(\boldsymbol{x},\boldsymbol{x}_n)$ representing the Euclidean distance between $\boldsymbol{x}$ and $\boldsymbol{x}_n$. $\boldsymbol{N}_{\mathrm{P}}$ and $\boldsymbol{N}_{\mathrm{D}}$ are AWGN noises. To exploit (2) in the estimation, the dependency on nuisance parameter $\boldsymbol{s}_{\mathrm{D}}$ must be eliminated. Two distinct approaches are proposed.

III. Philosophy and Derived Estimators

1) Projection-based estimator: Without any a priori information on $\boldsymbol{x}$—thus in a Fisher context—the localization problem can be formulated as a Maximum Likelihood (ML) estimation:

$$\widehat{\boldsymbol{x}} = \underset{\widetilde{\boldsymbol{x}}}{\operatorname{argmin}} \min_{\widetilde{\boldsymbol{s}}_{\mathrm{D}}\in\mathcal{C}_{\mathrm{D}}^{D}} ||\boldsymbol{Y}_{\mathrm{P}} - \boldsymbol{a}(\widetilde{\boldsymbol{x}})\boldsymbol{s}_{\mathrm{P}}^{T}||^2 + ||\boldsymbol{Y}_{\mathrm{D}} - \boldsymbol{a}(\widetilde{\boldsymbol{x}})\widetilde{\boldsymbol{s}}_{\mathrm{D}}^{T}||^2. \tag{4}$$

By relaxing the constraint on data constellation with a minimization on $\mathbb{C}^{D}$ instead of $\mathcal{C}_{\mathrm{D}}^{D}$, an estimation of $\boldsymbol{s}_{\mathrm{D}}$ can be constructed as $\widehat{\boldsymbol{s}}_{\mathrm{D}}^{T}(\widetilde{\boldsymbol{x}}) = \boldsymbol{a}^{\dagger}(\widetilde{\boldsymbol{x}})\boldsymbol{Y}_{\mathrm{D}}$ and injected back into (4), yielding a projection $\boldsymbol{P}(\widetilde{\boldsymbol{x}}) = \boldsymbol{a}(\widetilde{\boldsymbol{x}})\boldsymbol{a}^{\dagger}(\widetilde{\boldsymbol{x}})$. This approach leads to our **projection**-based estimator, which ignores the symbol structure but benefits from the model's knowledge in (2).

2) Transformation-based estimators: Still within the Fisher framework, another approach to leverage data observations and eliminate the $\boldsymbol{s}_{\mathrm{D}}$ dependency stems from the fact that, although unknown, the received data symbols across all nodes originate from the same transmitted symbols. By splitting the N antennas into two groups—of N_A and N_B antennas respectively—observations in group B can be expressed relatively to those in group A through a transformation $\boldsymbol{a}_B(\widetilde{\boldsymbol{x}}) = \boldsymbol{T}_{AB}(\widetilde{\boldsymbol{x}})\boldsymbol{a}_A(\widetilde{\boldsymbol{x}})$. This follows an extended time difference of arrival (TDOA) approach:

$$\boldsymbol{Y}_{\mathrm{D},B} = \boldsymbol{T}_{AB}(\widetilde{\boldsymbol{x}})\boldsymbol{Y}_{\mathrm{D},A} + \boldsymbol{N}_{\mathrm{D},AB}(\widetilde{\boldsymbol{x}}) \quad \in \mathbb{C}^{N_B\times D}, \tag{5}$$

where observations $\boldsymbol{Y}_{\mathrm{D},A}$ are thus considered as pilots and $\boldsymbol{N}_{\mathrm{D},AB}(\widetilde{\boldsymbol{x}})$ is the resulting noise. In this **relative** mode, the likelihood of all observations reveals a final factor $P[\boldsymbol{Y}_{\mathrm{D},A};\widetilde{\boldsymbol{x}}]$ which incorporates data constellation information. Leveraging these relative observations either leads to our **hybrid transformation-NDA** estimator (incorporating the last term), or the simpler **transformation**-based estimator (neglecting it).

IV. Results

The localization root mean square error (RMSE) is evaluated for these estimators, which are benchmarked against an estimator relying only on pilots.

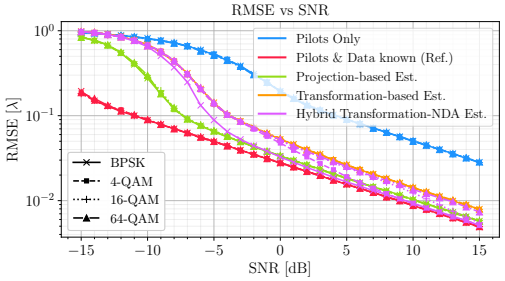

Fig. 1. RMSE compared to SNR with $P = 1$ (BPSK) $D = 32$, $N = 8$, $N_A = 5$ and $N_B = 3$. The receiver nodes form a circular array.

V. Extensions

Future work focuses on applying these approaches to more complex scenarios, incorporating random phase terms and attenuation, as well as exploiting the linear phase increase across subcarriers to improve localization capabilities.

References

[1] C. Mensing et al., "Data-Aided Location Estimation in Cellular OFDM Communications Systems," GLOBECOM 2009, 2009, pp. 1-7.

Poster: Through-Screen Computing

Hanting Ye
Delft University of Technology
Delft, The Netherlands
h.ye-1@tudelft.nl

Qing Wang
Delft University of Technology
Delft, The Netherlands
qing.wang@tudelft.nl

ABSTRACT

Mobile devices are playing increasingly significant roles in our daily lives through innovative mobile computing applications. Meanwhile, we observed the rise of *transparent screens* and their novel applications in advanced full-screen devices, whose front-facing optical sensors, such as ambient light sensors and cameras, are now placed under the transparent screen to capture ambient light and visual information. This design eliminates the on-screen area occupied by optical sensors, maximizing devices' screen-to-body ratio for the best use experience and device aesthetics. Motivated by this trend, we propose **Through-Screen Computing**, a new concept that we define as: *the computing of light signals for various purposes such as communication, sensing, and imaging, where the light comes from the physical world and passes through a special medium –the transparent screen– before reaching the under-screen optical sensors.* In this paper, we present the main challenges brought by transparent screens with respect to the proposed computing behind transparent screens. We describe how to overcome these challenges to retain the full functionality of under-screen sensors in terms of imaging and connectivity. Besides, we discuss some applications that could be enabled/enhanced by through-screen computing.

KEYWORDS

Through-screen computing, transparent screen, under-screen sensors, full-screen devices

1 INTRODUCTION

Mobile devices, such as smartphones, tablets, laptops, smartwatches, e-readers, and handheld gaming consoles, have become ubiquitous worldwide. Now, it is hard to imagine a world without these mobile devices. By 2022, there were more smartphones than people worldwide, and the number of mobile devices continues to grow at nearly five times the rate of the global human population [5]. The academic and industrial communities envision an exciting mobile computing future where mobile devices play an increasingly significant role in daily life. The rapid evolution in mobile devices leads to the fact that a flagship device from just a few years ago now seems outdated. Take the smartphone screens for an example. The screen has become the 'only' interactive interface between users and their smartphones since the launch of the first iPhone in 2007, which revolutionized the industry by eliminating most of the physical buttons. In the past decade, various innovative smartphone screen designs have been further realized to minimize the bezels and the notch area taken up by front cameras and other optical sensors to increase the screen-to-body ratio. These designs include the notch screen, teardrop notch screen, and through-hole screen of Android phones, as illustrated in Figure 1, and the "dynamic island" of iPhones. The ultimate goal is to achieve a borderless *"full-screen device"*, unifying user interaction functions and aesthetic design with the potential wide adoption of *transparent screens*.

Figure 1: Evolution of mobile devices and the efforts made on the screens to eliminate notch and bezel.

1.1 The Rise of Transparent Screen

Transparent screen technology utilizes transparent electrode materials to maintain display functionality while being visually transparent. Nowadays, transparent screens have revolutionized mobile devices, leading to the development of full-screen devices such as laptops (e.g., Thunderobot T-BOOK and Samsung Blade Bezel) and smartphones (e.g., ZTE AXON20/30/40, Xiaomi MIX4, and Samsung Galaxy Z Fold3/4/5/6) [3]. These full-screen devices, with their larger screen-to-body ratios, provide a better user experience by offering a more immersive and intelligent interface [1].

To achieve this goal, the screen of full-screen devices comprises a *Transparent Screen Region* and a *Normal Screen Region*, as illustrated in Figure 2(a-b). The transparent screen region is built with transparent electrode materials, serving two purposes: 1) displaying various contents, similar to the normal screen, and 2) allowing light to pass through the screen to reach under-screen optical sensors. The transparent screen's pixel layout is therefore optimized to balance the display functionality and the light transmittance to meet these two purposes [6], as shown in Figure 2(c). This innovative design allows placing optical sensors under the transparent screen without sacrificing their functionality, leading to the so-called *Under-Screen Sensors* [4], such as the Under-Screen Ambient Light Sensors (ALS) [1] and Under-Screen Cameras (USC)[1].

1.2 Our Vision: Through-Screen Computing

Transparent screens are revolutionizing our visual experience of mobile devices. However, they also change the traditional mobile computing of light-based signals since optical sensors now must be placed *under* the transparent screen instead of traditionally *on* the screen. Motivated by this paradigm shift, in this paper, we propose the concept of **Through-Screen Computing**, which we define

[1]It is also referred as Under-Display Camera (UDC) or Under-Panel Camera (UPC) in the literature [2, 6–8].

Hanting Ye and Qing Wang

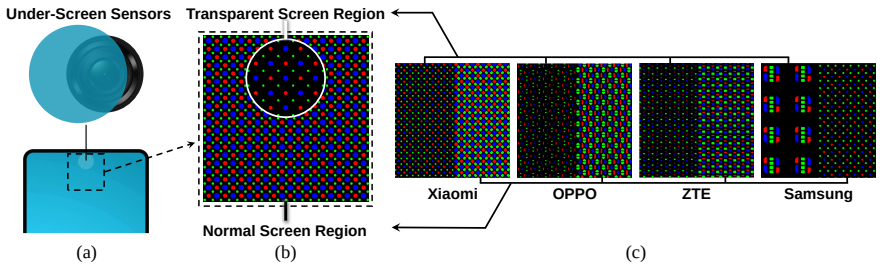

Figure 2: Illustrations: (a) full-screen smartphone with under-screen sensors; (b) magnified micrograph of the transparent screen region and the normal screen region; (c) screen diversity: the comparison of transparent screen regions and normal screen regions designed by different smartphone manufacturers.

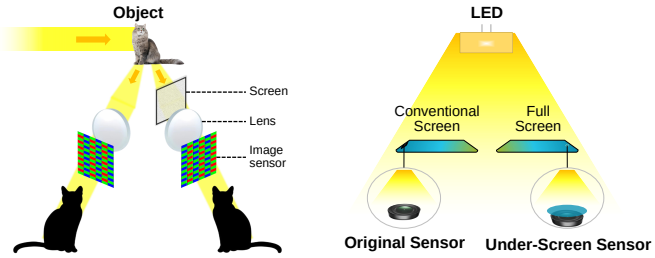

Figure 3: The concept of Through-Screen Computing: (left) Passive source; (right) Active source.

as: *the computing of light signals for various purposes such as communication, sensing, and imaging, where the light comes from the physical world and passes through a special medium –the transparent screen– before reaching the under-screen optical sensors.* We will explore several key themes with long-term potential for research and innovation: (i) How to overcome screen barriers for existing important directions in mobile computing, such as imaging and communication; (ii) How to leverage the screen to drive innovation in mobile computing, such as interaction, sensing, privacy, security, energy harvesting, and in fields like Augmented/Mixed Reality (AR/MR) and spatial intelligence. Below, we first present the essential components of the envisioned through-screen computing.

Light Source: Through-screen computing uses light signals as the computing input. We mainly consider two types of light sources: (1) *Passive sources:* referring to the objects that reflect light when illuminated, containing spatial information about the object and its environment. (2) *Active sources:* referring to the LED luminaires present in the lighting infrastructure, which are not only beneficial for illumination but also can be modulated in light intensity or colors at a high frequency to transmit information.

Transparent Screen: In through-screen computing, the transparent screen significantly affects light propagation. It influences the visual imaging of light reflected by passive objects (e.g., cats, see Figure 3(left)) or the intensity and color of light emitted by active sources (e.g., LEDs, see Figure 3(right)). Additionally, the transparent screen can act as an active source, featuring an array of RGB pixels in various layouts, shapes, and sizes. This array displays dynamic contents on mobile devices and brings challenges to through-screen computing, such as attenuation, color shift, and interference with other passive and active light.

Under-Screen Sensors: We consider two types of under-screen sensors in this paper: (1) *Single-pixel under-screen sensors*, such as an under-screen photodiode and an under-screen ambient light sensor. Both are semiconductor devices that convert the detected light into electrical information. While an under-screen photodiode only measures light intensity, an under-screen ambient light sensor, which combines a photodiode and a color filter, can measure both the intensity and the color of through-screen light signals; and (2) *Multi-pixel under-screen sensors*, such as an under-screen camera, which can also detect through-screen light signals but in an intuitive and understandable multi-pixel image output.

In this vision paper, we will focus on the computing behind the transparent screens of full-screen devices, addressing several critical challenges to advancing through-screen computing.

REFERENCES

[1] ams. 2019. *ams launches optical sensor which measures ambient light from behind a smartphone's OLED screen.*
[2] Neil Emerton, David Ren, and Tim Large. 2020. 28-1: Image Capture Through TFT Arrays. In *SID Int. Symp., Dig. Tech. Pap.*
[3] P.Gagnon. 2019. Presentation at OLED World Summit 2019. *AMOLED Market & Technology Trend* (2019).
[4] Samsung. 2022. *What is the Under Display Camera on the Galaxy Z Fold 4?*
[5] Statista. 2023. *Charted: There are more mobile phones than people in the world.*
[6] Zhibin Wang and et.al. 2020. Self-Assembled Cathode Patterning in AMOLED for Under-Display Camera. In *SID Int. Symp., Dig. Tech. Pap.*
[7] Anqi Yang and Aswin C Sankaranarayanan. 2021. Designing display pixel layouts for under-panel cameras. *IEEE TPAMI* (2021).
[8] Yuqian Zhou, David Ren, Neil Emerton, Sehoon Lim, and Timothy Large. 2021. Image restoration for under-display camera. In *CVPR.*

The authors of this abstract did not consent to publishing their abstract in the SITB proceedings

On the Privacy-Robustness Trade-off in Distributed Average Consensus

Zarè Palanciyan, Delft University of Technology, z.palanciyan@student.tudelft.nl,
Qiongxiu Li, Aalborg University, qili@es.aau.dk,
Richard Heusdens, Netherlands Defence Academy, Delft University of Technology, r.heusdens@tudelft.nl

***Abstract*—Distributed consensus algorithms face a dual challenge in modern networked systems: safeguarding sensitive data through privacy-preserving mechanisms while maintaining robustness against adversarial nodes (e.g. Byzantine faults). This paper shows a fundamental trade-off between privacy preservation and adversarially detection in distributed average consensus algorithms. Our analysis reveals that as the level of privacy preservation increases, the detection accuracy for adversarials declines accordingly, highlighting a critical design challenge for secure and privacy-aware consensus frameworks.**

***Index Terms*—ADMM, privacy, subspace perturbation, adversary, detection, median absolute deviation, privacy-robustness trade-off**

I. Introduction

Maintaining robustness against adversarial nodes is a necessary goal in distributed computations. Without this robustness, the output of distributed computations will fall into the hands of adversaries. This will render distributed computations useless, as they will not be able to converge to their optimal values. But as long as an adversarial detection algorithm can identify whether a node is honest or adversarial based on its outputs, it will guarantee that the distributed computation will achieve its optimal value. However, as we will show, having full knowledge of a node raises privacy issues. As distributed computations are applied across more fields, they increasingly handle sensitive data, and privacy protection is vital [1]. Meaning that having full knowledge of a node is not optimal in some cases, this introduces a trade-off between privacy and robustness in distributed computations.

II. Methods

To show this trade-off in effect, we used the ADMM algorithm for distributed average consensus computation, where we implemented the subspace perturbation (SP) privacy-preserving mechanism. This mechanism inserts noise into the non-convergent subspace of the auxiliary variables of the ADMM algorithm, which does not affect output accuracy while being able to achieve perfect privacy [2]. For robust detection, we slide a window of L iterations and compute the median absolute deviation (MAD) of all incoming updates on each edge, following [3]. If neighbour j surpasses the threshold more than $L/2$ times within the window, node i marks j as adversarial and ignores its messages for the subsequent L iterations. We assume that the adversary will perform a Gaussian noise attack where the ADMM updates are replaced by random noise, making the consensus average to deviate from the true average.

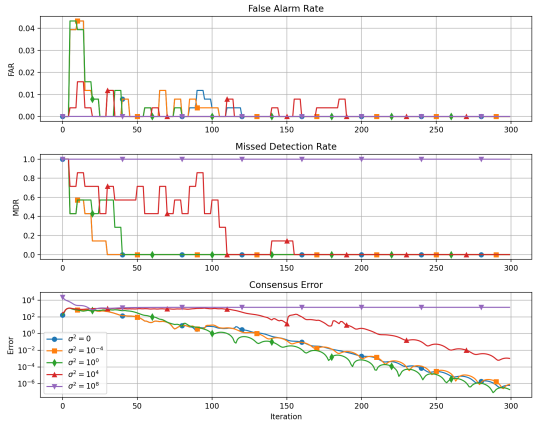

Fig. 1. ADMM simulation under privacy-preserving subspace perturbation with noise variances $\sigma^2 \in \{0, 10^{-4}, 1, 10^4, 10^8\}$. From top to bottom: (i) *False Alarm Rate (FAR)*: fraction of honest edges incorrectly flagged; (ii) *Missed Detection Rate (MDR)*: fraction of truly corrupt edges that go undetected; (iii) *Consensus Error* $\|x_i - x^\star\|^2$ on a log scale.

III. Numerical Results And Conclusion

Figure 1 shows the trade-off between privacy and robustness. A higher inserted SP noise variance will raise the privacy and make detection less likely. Hence, as the noise variance grows, the corrupt nodes become more indistinguishable from honest nodes. These results expose a fundamental trade-off between privacy and robustness in consensus protocols: the stronger the privacy guarantee, the harder it becomes to identify Byzantine behaviour.

References

[1] M. Mageshwari and R. Naresh, "Decentralized data privacy protection and cloud auditing security management," in *2022 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)*, 2022, pp. 103–109.

[2] Q. Li, J. S. Gundersen, M. Lopuhaä-Zwakenberg, and R. Heusdens, "Adaptive differentially quantized subspace perturbation (adqsp): A unified framework for privacy-preserving distributed average consensus," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 1780–1793, 2024.

[3] O. Shalom, A. Leshem, and A. Scaglione, "Localization of data injection attacks on distributed m-estimation," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 8, pp. 655–669, 2022.

Machine Learning-based Lifetime Prediction and Uncertainty Estimation of GaN HEMTs

Shuoyan Zhao, Raj Thilak Rajan, Andrea Natale Tallarico, Maurizio Millesimo, Vladislav Volosov, Antonio Imbruglia, and Justin Dauwels

Abstract

Accurate prediction of Gallium Nitride High-Electron Mobility Transistors (GaN HEMTs) lifetime is essential for ensuring the reliability of power electronics [1]. Traditional model-based approaches offer physically interpretable predictions but often struggle to account for stochastic variability observed in real-world scenarios. Conversely, machine learning (ML) methods perform well in complex environments but often fail in low-resource and out-of-distribution settings [2]. Additionally, limited incorporation of domain knowledge reduces interpretability—particularly regarding uncertainty quantification-and hinders their acceptance in industrial deployments.

This work proposes an ML-based framework for static lifetime prediction of GaN HEMTs, leveraging early-stage degradation signals with an emphasis on uncertainty estimation. The University of Bologna tested 49 devices [3] under varying gate voltages and temperatures. Time-series measurements of gate current, drain current, and ON-resistance were recorded every two seconds until device failure. A log transformation is applied to lifetime values to normalize their distribution and reduce skew. Additionally, time-domain features were extracted from the first two seconds of measurements to capture early characteristics. We investigate two algorithms: (i) ensembles of bootstrapped XGBoost regressor [4] for empirical uncertainty estimation, and (ii) Gaussian Process (GP) regressor that inherently captures predictive uncertainty. Additionally, we introduce monotonicity constraints on selected features—such as temperature and gate voltage—to incorporate domain priors and assess their effect on reducing epistemic uncertainty without fully specifying a physical model.

Results, validated under leave-one-device-out cross-validation (LOOCV), indicate that XGBoost achieves best accuracy (log-SMAPE = 11.6%), while GP models provide better-calibrated uncertainty estimates, as shown in Fig. 1. Though not ideal, the models demonstrate promising capability for approximate lifetime estimation in industrial applications, offering valuable insights for further development. However, the current approach to enforcing monotonicity in GP models relies on additive kernel assumptions [5], which imply feature independence—contradicting observed interdependencies and distorting predictions. Future work will explore more expressive constraint formulations (e.g., local boundary conditions) and expand the dataset to improve model generalization and interpretability.

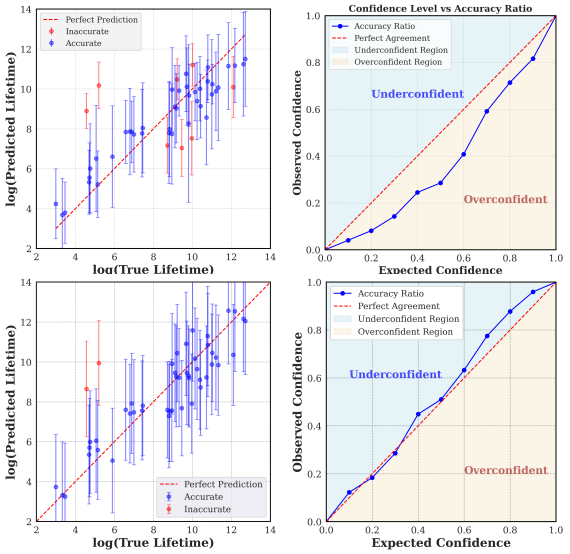

Fig. 1: Comparison of prediction accuracy and uncertainty calibration for XGBoost (top) and Gaussian Process (bottom) under LOOCV. Left: Predicted vs. true log-lifetime with 95% confidence intervals; Right: Calibration curves showing observed vs. expected confidence and miscalibration regions.

Acknowledgments

This work, as a part of R-PODID project, is supported by the Chips Joint Undertaking and its members, including the top-up funding by National Authorities of Italy, Turkey, Portugal, The Netherlands, Czech Republic, Latvia, Greece, and Romania under grant agreement n° 101112338.

References

[1] J. P. Kozak, R. Zhang, M. Porter, Q. Song, J. Liu, B. Wang, R. Wang, W. Saito, and Y. Zhang, "Stability, reliability, and robustness of gan power devices: A review," *IEEE Transactions on Power Electronics*, vol. 38, no. 7, pp. 8442–8471, 2023.

[2] H. Li, Z. Zhang, T. Li, and X. Si, "A review on physics-informed data-driven remaining useful life prediction: Challenges and opportunities," *Mechanical Systems and Signal Processing*, vol. 209, p. 111120, 2024.

[3] STMicroelectronics, "Sgt65r65al - 650 v, 65 mω, aec-q101 power gan transistor," 2025. Accessed: 2025-04-14.

[4] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

[5] A. F. López-Lopera, F. Bachoc, N. Durrande, and O. Roustant, "Finite-dimensional gaussian approximation with linear inequality constraints," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 6, no. 3, pp. 1224–1255, 2018.

S.Zhao, R.T.Rajan and J.Dauwels are with the Signal Processing Systems Group, Delft University of Technology, Delft, The Netherlands (e-mail: S.Zhao-2@tudelft.nl, R.T.Rajan@tudelft.nl, J.H.G.Dauwels@tudelft.nl).

A.N.Tallarico, M.Millesimo, V.Volosov are with the Advanced Research Center on Electronic System, Department of Electrical, Electronic, and Information Engineering, University of Bologna, Cesena, Italy (e-mail: a.tallarico@unibo.it, maurizio.millesimo2@unibo.it, vladislav.volosov2@unibo.it).

A.Imbruglia is is with STMicroelectronics, Catania, Italy (e-mail: antonio.imbruglia@st.com).