

Learned layered coding for Successive Refinement in the Wyner-Ziv Problem

Joukovsky, Boris; De Weerd, Brent; Deligiannis, Nikos

Published in:

2024 IEEE International Conference on Acoustics, Speech and Signal Processing

DOI:

[10.1109/ICASSP48485.2024.10446574](https://doi.org/10.1109/ICASSP48485.2024.10446574)

Publication date:

2024

Document Version:

Accepted author manuscript

[Link to publication](#)

Citation for published version (APA):

Joukovsky, B., De Weerd, B., & Deligiannis, N. (2024). Learned layered coding for Successive Refinement in the Wyner-Ziv Problem. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 6020-6024). (ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings). IEEE. <https://doi.org/10.1109/ICASSP48485.2024.10446574>


Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.


Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

LEARNED LAYERED CODING FOR SUCCESSIVE REFINEMENT IN THE WYNER-ZIV PROBLEM

Boris Joukovsky^{1,2} 

Brent De Weerd^{1,2} 

Nikos Deligiannis^{1,2} 

¹ETRO Department, Vrije Universiteit Brussel (VUB), Pleinlaan 2, B-1050 Brussels, Belgium
²imec, Kapeldreef 75, B-3001 Leuven, Belgium

ABSTRACT

We propose a data-driven approach to explicitly learn the progressive encoding of a continuous source, which is successively decoded with increasing levels of quality and with the aid of correlated side information. This setup refers to the successive refinement of the Wyner-Ziv coding problem. Assuming ideal Slepian-Wolf coding, our approach employs recurrent neural networks (RNNs) to learn layered encoders and decoders for the quadratic Gaussian case. The models are trained by minimizing a variational bound on the rate-distortion function of the successively refined Wyner-Ziv coding problem. We demonstrate that RNNs can explicitly retrieve layered binning solutions akin to scalable nested quantization. Moreover, the rate-distortion performance of the scheme is on par with the corresponding monolithic Wyner-Ziv coding approach and is close to the rate-distortion bound.

Index Terms— Wyner-Ziv coding, successive refinement, layered coding, nested scalar quantization, recurrent neural networks

1. INTRODUCTION

Distributed source coding (DSC) considers the separate encoding and joint decoding of samples of correlated sources. In the asymmetric scenario where one discrete source Y is encoded independently, the Slepian-Wolf (SW) theorem [1] implies that a correlated source X can be decoded losslessly using Y as side information, while being compressed at a rate identical to the case where the two sources are encoded jointly. Notably, this result has been generalized to Wyner-Ziv Coding (WZC) in the case of jointly Gaussian sources with square distortion measure [2], extending SW coding to lossy compression. *Successive refinement* for WZC is defined as the progressive, quality scalable coding of the source with available side information at the decoder. Crucially, successive refinability is achieved when no rate loss occurs compared to monolithic (non-scalable) coding, for instance, in the case of the joint Gaussian source and non-degraded side information [3].

Designing practical codes for WZC remains a challenge, in part because the achievability of SW coding (SWC) relies on a non-constructive, random binning procedure. Practical solutions approaching the rate-distortion bound include the use of trellis codes [4] or nested lattice codes [5] for quantization (sphere covering) and channel codes such as low density parity-check (LDPC) codes for SWC (sphere packing). The side information is used to decode both the channel and source codes. The approaches [6, 7] to design successive refinable codes employ nested scalar quantization (NSQ) followed by SW encoding and decoding of successive bit-planes of the quantized signal using LDPC codes. However, NSQ requires very high dimensions to reach optimality and decoding

LDPC codes is computationally expensive due to the use of belief propagation on a factor graph.

More recently, neural compression has been a subject of research in deep learning. The study in [8] has shown that black-box deep variational auto-encoders can learn a compressed representation of images, by optimizing a rate-distortion trade-off. Following these ideas, resolution scalable compression of images has also been studied in the non-distributed case [9], as well as quality scalable approaches [10]. Only a few works have recently attempted to design neural compression schemes for WZC, with applications in distributed image compression [11, 12], multi-view image coding [13], and federated learning topics, that is, distributed compression of gradients [14] and federated compression learning [15]. Despite their success, these existing neural WZC methods use heuristic approaches, where an autoencoder is used for separate encoding and joint decoding of correlated signals; these approaches are therefore not interpretable in the information theoretic sense of WZC. Recent progress towards understanding learned WZ codes has been made in [16], which showed that artificial neural networks (ANNs) trained to minimize an upper bound on mutual information can explicitly learn random binning structures, under the assumption of ideal SWC or ideal entropy coding. This is a first-of-its-kind result in data-driven approaches because it is akin to the achievability argument of the WZ theorem for jointly Gaussian sources.

In this paper, we present the first attempt to learn successive refinement for the Wyner-Ziv problem by extending [16] to successively refined random binning. Specifically, we show that recurrent neural networks (RNNs) are suitable candidates to efficiently learn refinable encoders and decoders, akin to achieving layered WZC. Unlike vector quantized-based architectures [11], we rather let the model learn explicit quantizers that exhibit a behavior akin to binning of NSQ planes, which is known to be an asymptotically optimal solution to successively refinable WZC [6]. Our approach is based on optimizing two variational bounds on the asymptotic rates of each encoder, which are the rates of an ideal entropy coder and an ideal SW encoder, respectively. We show that the second case approaches the theoretical lower bound of layered WZC, while achieving equal performance to the monolithic scheme of [16] and retrieving similar solutions to bit-plane binning.

The paper is organized as follows: Section 2 describes the background on the successive refinement of the WZ Problem. Section 3 derives the variational loss functions used to train the refinement stages and the proposed RNN-based architecture. Section 4 presents the results obtained by training the model on correlated sources—whose correlation is modeled by an additive i.i.d. Gaussian noise—and explicitly displays the learned binning structure. Lastly, Section 5 concludes the work.

2. BACKGROUND

2.1. Wyner-Ziv coding

We start by reminding the Wyner-Ziv theorem for DSC [2]. Consider a discrete, memoryless source $X \in \mathcal{X}^n$ and correlated side information $Y \in \mathcal{Y}^n$ known at the decoder-side only. The decoded signal $\hat{X} \in \hat{\mathcal{X}}^n$ is reconstructed with an average distortion D based on the distance function $d(X, \hat{X})$, using Y and a compressed bit string at some rate R obtained by separate encoding of X . The rate-distortion function $R_{\text{WZ}}(D)$ defining the set of achievable rate-distortion pairs in this setup is given by:

$$R_{\text{WZ}}(D) = \min_{\substack{p(u|x) \\ \mathbb{E}d(X, \hat{X}) \leq D}} I(X; U|Y), \quad (1)$$

where U is an auxiliary variable satisfying the Markov chains $U \leftrightarrow X \leftrightarrow Y$ and $X \leftrightarrow (U, Y) \leftrightarrow \hat{X}$. Remarkably, under certain conditions [2, 17, 18], $R_{\text{WZ}}(D)$ achieves the same rate-distortion performance as when the side information is known also at the encoder.

2.2. Successively refined Wyner-Ziv coding

A successive refinement code with K stages and non-degraded side information is defined as a set of encoder-decoder pairs (ϕ_k, ψ_k) , where the decoders successively estimate \hat{X}_k across distortion levels (D_1, \dots, D_K) with the aid of the side information and previously accumulated codes:

$$\begin{aligned} \phi_k : \mathcal{X}^n &\rightarrow \{1, 2, \dots, M_k\}, \\ \psi_k : \{1, \dots, M_1\} \times \dots \times \{1, \dots, M_k\} \times \mathcal{Y}^n &\rightarrow \hat{\mathcal{X}}^n, \end{aligned} \quad (2)$$

such that $\mathbb{E}d(X^n, \psi_k(\phi_1(X^n), \dots, \phi_k(X^n), Y^n)) \leq D_k$ and M_k are alphabet sizes. The individual encoders communicate at rates (R_1, \dots, R_K) and the total *sum-rate* sent after k stages is given by $R_k^{\text{tot}} = \sum_{i=1}^k R_i$. The conditions for successive refinability in the Wyner-Ziv setting are stated next:

Theorem [3]: X is said to be successively refinable with distortion levels (D_1, D_2, \dots, D_K) , if and only if there exists auxiliary random variables U_k for $k = 1, \dots, K$, with K and the deterministic functions $f_k : \mathcal{U}_k \times \mathcal{Y} \rightarrow \hat{\mathcal{X}}$ such that the three following conditions hold, $\forall k$:

1. $R_{\text{WZ}}(D_k) = I(X; U_k|Y)$ s.t. $\mathbb{E}[d(X, f_k(U_k, Y))] \leq D_k$,
2. $(U_1, \dots, U_K) \leftrightarrow X \leftrightarrow Y$,
3. $(U_1, \dots, U_{k-1}) \leftrightarrow (U_k, Y) \leftrightarrow X$, $k > 1$,

where $R_{\text{WZ}}(D)$ is the Wyner-Ziv rate-distortion function. In other words, each stage must achieve the same optimal rate (in terms of sum-rate) as the monolithic (non successively coded) scheme. We may also note that the differential rate-distortion function, which characterizes the rate of each refinement layer, is given by:

$$R_{\text{WZ}}(D_k) - R_{\text{WZ}}(D_{k-1}) = I(X; U_k|U_{k-1}, \dots, U_1, Y) \quad (4)$$

The study in [3] proved that WZC is successively refinable when X and Y are joint Gaussian and the ℓ_2 is the distortion metric. [6] relaxed the condition to random sources correlated by additive *i.i.d.* Gaussian noise.

3. LEARNING SUCCESSIVELY REFINED WYNER-ZIV

We aim to jointly learn the quantization, binning and reconstruction of a successively refinable source with side information at the decoder, with asymptotic rates close to the ideal WZ rate. In the

way of [16], we seek to discover if a trained network can recover the NSQ and bit-plane binning solution, which is known to achieve successive refinability in WZC [6]. As for notations, the message resulting from the encoding of $x \sim X$ is denoted as $[m_1, \dots, m_K]$ where $m_k \in \{0, \dots, M-1\}$. The successive symbols m_k are the outputs of K encoders and are sent to K corresponding decoders. We abbreviate $[m_1, \dots, m_k]$ with m_1^k and likewise for sequences of any other variable. By abuse of notation, m_1^k for $k \leq 0$ refers to an empty sequence. Unless specified otherwise, $\log(\cdot)$ is the base-2 logarithm. In the following, we assume that the side information Y is a noisy version of the source, such that $X = Y + N$, where $N \sim \mathcal{N}(0, \sigma_n^2)$ is Gaussian with different possible noise levels. This scenario is proven to be successively refinable in the WZ setting [6].

3.1. Objective function

We first discuss the training objective for a learned successive refinement model. For each refinement stage k , with reconstruction outputs \hat{x}_k , we use as distortion metric $\mathbb{E}[d(x, \hat{x}_k)]$, where the expectation operator averages over all possible inputs x and reconstructions \hat{x}_k , and the distance metric d is the mean squared error (MSE). Recall then from Eq. (4) that the differential rate of a single stage R_k for the ideal WZ coder is given by:

$$R(D_k) = \min I(X; U_k|U_1^{k-1}, Y) \quad (5)$$

$$= \min H(U_k|U_1^{k-1}, Y) - H(U_k|U_1^{k-1}, X) \quad (6)$$

$$= \mathbb{E}_{\substack{x, y \sim p(x, y) \\ u_1^k \sim p(u_1^k|x)}} \left[\log \frac{p(u_k|u_1^{k-1}, x)}{p(u_k|u_1^{k-1}, y)} \right], \quad (7)$$

with the minimization constrained by $\mathbb{E}[d(x, \hat{x}_k)] \leq D_k$. When the coding problem is successively refinable, each stage k asymptotically achieves the WZ rate-distortion bound.


Since the distribution $p(u_k|u_1^{k-1}, y)$ in general is not known in closed form, we define two upper bounds for (7), similar to [16]:

$$R(D_k) \leq \mathbb{E}_{\substack{x, y \sim p(x, y) \\ u_1^k \sim p(u_1^k|x)}} \left[\log \frac{p(u_k|u_1^{k-1}, x)}{q_m(u_k|u_1^{k-1})} \right], \quad (8)$$


$$R(D_k) \leq \mathbb{E}_{\substack{x, y \sim p(x, y) \\ u_1^k \sim p(u_1^k|x)}} \left[\log \frac{p(u_k|u_1^{k-1}, x)}{q_c(u_k|u_1^{k-1}, y)} \right], \quad (9)$$

where q_m and q_c are two different models for the distribution $p(u_k|u_1^{k-1}, y)$ and referred to as *prior models*. These two bounds can be interpreted as corresponding to two different coding systems. The bound in (8), using the marginal distribution $q_m(u_k|u_1^{k-1})$, is the rate of a system where the output of each encoder stage $p(u_k|u_1^{k-1}, y)$ is further compressed using an ideal entropy coder at the rate $\mathbb{E}_x \left[\mathbb{E}_{u_1^k \sim p(u_1^k|x)} [-\log q_m(u_k|u_1^{k-1})] \right]$. Eq. (9) uses the conditional distribution $q_c(u_k|u_1^{k-1}, y)$, where the output of the encoder is compressed using an ideal Slepian-Wolf coder, asymptotically reaching a rate of $\mathbb{E}_{x, y} \left[\mathbb{E}_{u_1^k \sim p(u_1^k|x)} [-\log q_c(u_k|u_1^{k-1}, y)] \right]$.

To establish the connection with the practical refinement code, we consider that each sample from U_k in Theorem 1 represents the discrete encoder messages $[m_1, \dots, m_k]$, such that $p(u_k) = p(m_k^k)$ and $p(u_k|u_1^{k-1}) = p(m_k|m_1^{k-1})$. We also relax the hard constraint $\mathbb{E}[d(x, \hat{x}_k)] \leq D_k$ in the WZ theorem to obtain one of the two following losses for a single stage, depending on the choice of the



(a) Proposed RNN-based encoders and decoders.



(b) Marginal and conditional RNN-based prior modules.

Fig. 1: Our successive refinement model. Gray boxes are RNNs.

upper bounds between (8) or (9):

$$\mathcal{L}_{m,k} = \mathbb{E} \left[\log \frac{p(m_k | m_1^{k-1}, x)}{q_m(m_k | m_1^{k-1})} + \lambda d(x, g_k(m_k, y)) \right], \quad (10)$$

$$\mathcal{L}_{c,k} = \mathbb{E} \left[\log \frac{p(m_k | m_1^{k-1}, x)}{q_c(m_k | m_1^{k-1}, y)} + \lambda d(x, g_k(m_k, y)) \right], \quad (11)$$

where $g_k(\cdot)$ is a joint decoder at stage k with a continuous output, while the encoders and the prior models will be learned using discrete probabilistic models with categorical outputs. The derived loss functions minimize the rate and distortion simultaneously, balanced by hyperparameter λ , and the expectation operator subscripts are the same as in (8)-(9), but are left out for brevity. In practice, the complete, K -stages model is trained globally and end-to-end by summing the losses for each stage into the total loss function $\mathcal{L}_m = \sum_k \mathcal{L}_{m,k}$ or $\mathcal{L}_c = \sum_k \mathcal{L}_{c,k}$. The expectations will be approximated by training the models using stochastic gradient descent on large batches, and by sampling data samples x, y from the input distribution and $[m_1, \dots, m_K]$ using the encoding functions $p(m_k | m_1^{k-1}, x)$. In the following section, we discuss the models used to learn the encoder and prior distributions, as well as the decoding functions.

3.2. Proposed RNN-based successive refinement model

Our model is shown in Fig. 1 and relies on stacked RNNs whose hidden states encode information of lower refinement levels. As depicted in Fig. 1a, the K encoders $p(m_k | m_1^{k-1}, x)$ are parameterized by $\text{softmax}(f(h_k^{(L)}))$, where $h_k^{(L)}$ is the hidden state of a stacked RNN with L layers, taking the source realization x as input for every time-step. We share a linear layer f across all K refinement stages. Therefore, after normalization, each encoder yields the probability vector $[p(m_k = 0 | m_1^{k-1}, x), \dots, p(m_k = M | m_1^{k-1}, x)]$, which is used in the previously derived loss functions. The successive decoders $g_k(m_k, y)$ are parameterized by $g(\hat{h}_k^{(L)})$, where $\hat{h}_k^{(L)}$ is the hidden state of another stacked RNN taking y and the k^{th} encoder output as input for every stage, assuming that the hidden state encodes all the previously observed codes. The prior modules in Fig. 1b yield $q_m(m_k | m_1^{k-1})$ and $q_c(m_k | m_1^{k-1}, y)$ as $\text{softmax}(p_k(\hat{h}_k^{(L)}))$, where p_k are linear mappings and $\hat{h}_k^{(L)}$ is the output of yet another RNN taking at every stage k either m_{k-1} or

the pair (y, m_{k-1}) as input, respectively. Setting $K = 1$ reduces our model to the monolithic scheme in [16].

The motivation of using RNNs lies in the observation that the decision functions of a traditional NSQ—which is essential for layered WZC—are described by a quantization of successively finer fractional parts of the input and require modular arithmetics. In this context, we find that RNNs are suitable to learn a recursive function of the input that approximate the quantization of successive bit-planes of a NSQ, thereby allowing scalability at higher quality levels. Moreover, at the decoder, using a RNN with sequential inputs maintains a fixed input size (one-hot encoded bins) compared to using separate decoders with increasing complexity. Lastly, using a RNN diminishes the risk of training separate feed-forward NNs with a potential competition between the refinement stages.

During training, we use the same variational training approach as in [16] by sampling the output of the encoder and the prior modules according to a Gumbel-softmax or Concrete distribution [19,20] with temperature τ . This τ is decreased during training, thereby avoiding instability during the initial epochs and eventually making the samples m_k approach a categorical (hard) distribution. During inference, the encoder is no longer probabilistic and the codes m_k are set to $\text{argmax}(f(h_k))$, encoded as one-hot vectors. This allows to estimate the theoretical rate for each level of the ideal entropy coder or SWC respectively, by summing over N samples (x_n, y_n) drawn from $p(x, y)$:

$$R_{\text{marg},k} \approx \frac{1}{N} \sum_{n=1}^N -\log q_m(m_{n,k} | m_{n,1}^{k-1}), \quad (12)$$

$$R_{\text{cond},k} \approx \frac{1}{N} \sum_{n=1}^N -\log q_c(m_{n,k} | m_{n,1}^{k-1}, y_n), \quad (13)$$

where the probabilities in Eq. (12)-(13) are selected from the prior modules according to the predicted codes.

4. RESULTS

All RNNs use 2 hidden layers with 100 hidden units each and with LeakyReLU activations, and linear layers are added at the output of each RNN, which are (Gumbel-)softmax-activated for the probabilistic models. We also compare with the monolithic scheme, by implementing [16] and training it for different numbers of bins. All models, including the monolithic ones, are trained using the Adam optimizer with an initial learning rate of 10^{-3} for 180 epochs. The learning rate is decreased by a factor 0.3 every 80 epochs for the marginal distribution model in (8), and every 40 epochs for the conditional model in (9). The Gumbel-softmax temperature τ is decreased exponentially from 1.0 to 0.2 during training. We draw 2×10^5 (x, y) samples per epoch with a batch size of 10^3 , s.t. $Y \sim \mathcal{N}(0, 1)$ and $X = Y + N$ where $N \sim \mathcal{N}(0, \sigma_n^2)$. During testing, we estimate the rates and distortions over 10^7 samples. For stability purposes, we use a stop-gradient operation on the predicted codes at the input of the prior modules as it was found to improve the convergence of the model.

We consider two successive refinement scenarios, and consequently two versions of the RNN models. In one case, we use three refinement stages, where the output of the encoder at each level is binary ($M = 2$). This scenario is referred to as “222” The rate-distortion should thus be similar to the output of the monolithic model with 2, 4, and 8 bins respectively. The other scenario has two refinement stages with an alphabet size of $M = 4$, called the “44” model. These two stages should then match the performance




Fig. 2: Rate-distortion curves for the monolithic model and the proposed successive refinement model in the 222 and 44 setting, for both the marginal and conditional rate estimations.

of the monolithic model with 4 and 16 bins, respectively. Furthermore, both marginal and conditional distribution models resulting from the upper bounds (8)-(9) are tested in both scenarios. The rate-distortion curves for all successive refinement RNNs are shown in Fig. 2, along with the marginal and conditional monolithic models from [16], the Wyner-Ziv bound given by (1) and the rate-distortion bound for when no side information is available. We will discuss the results for the marginal and conditional models separately in the following sections.

4.1. Marginal models

As mentioned before, in this setup resulting from the bound (8), the encoder needs to learn both the quantization of the signal and the entropy coding of the resulting bins. The rate of the system is calculated using the cross-entropy approximation of Eq. (12). The same binning behavior is observed in both the marginal and conditional models, and illustrated for the conditional case in Fig. 3.

For all models, we observe that the output distribution at each stage or number of bins is nearly uniform, resulting in rates only very slightly below $\log N$, with N the number of bins. In the left plot of Fig. 2, we can see that the successive 222 coder performs nearly as well as the monolithic model for the noise level $\sigma_n^2 = 0.1$, with differences in distortion below 0.4 dB. For $\sigma_n^2 = 0.01$ in the right plot of Fig. 2, the first stage of the successive 222 model improves over the monolithic model with a reduction of 1.4 dB in distortion. The successive 44 model loses 0.7 and 2.4 dB in the case $\sigma_n^2 = 0.1$, while for $\sigma_n^2 = 0.01$ the model loses 1.3 dB at the first stage, but improves on the monolithic model at the second stage with 1.0 dB.

4.2. Conditional models

In this setup, resulting from the bound (9), the encoder can be seen as a quantizer, while the distribution $q_c(m_k|m_1^{k-1}, y)$ relates to the rate of an ideal SW coder and is estimated according to Eq. (13). The results for this scenario are therefore closer to the WZ bound than the marginal models, since a larger part of the code is assumed to be ideal. For $\sigma_n^2 = 0.1$ in Fig. 2, the curves for the monolithic and successive refinement models are nearly identical, showing that we can achieve refinement without rate loss. For the low-noise case $\sigma_n^2 = 0.01$ in Fig. 2, the layered coders even improve over the monolithic model, with a reduction in distortion of up to 1.7 dB.

Contrary to [16], we observe discontinuous binning of source samples, resembling the random binning in the Slepian-Wolf achievability theorem [1], in both the marginal and conditional models. An illustration is given in Fig. 3a for the conditional 222 model and in Fig. 3b for the 44 model. In the top row, the output of the encoder is plotted for each stage, in function of the input sample x . The map-

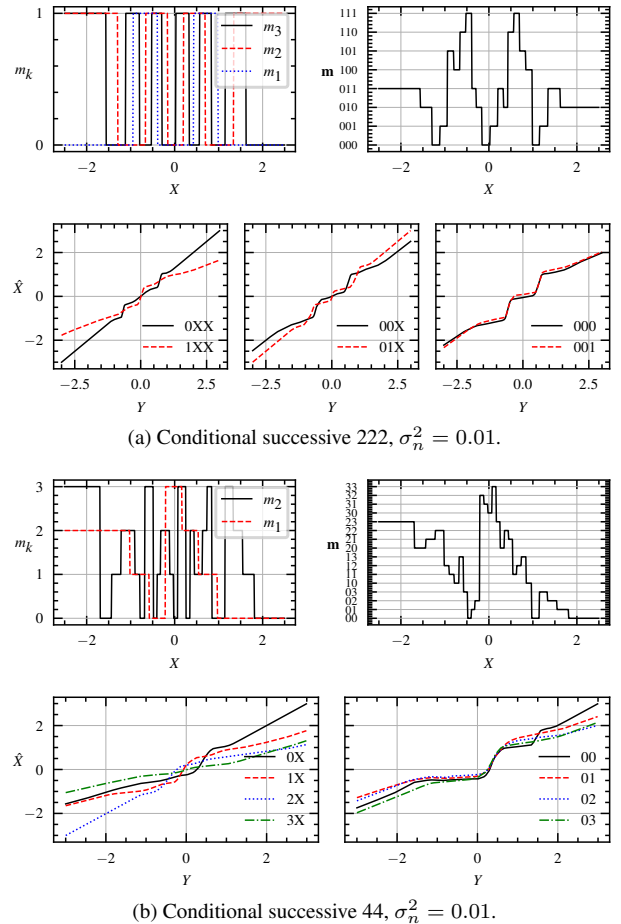


Fig. 3: Top-left: visualization of the learned quantizer or bins at each encoding layer. Top-right: complete message \mathbf{m} , by concatenation of the encoders outputs. Bottom: successively refined input reconstruction w.r.t the side information and sent codes (e.g.: '0XX' shows the output of the first decoding stage with $m_1 = 0$)

pings learned by the RNNs are characterized by discontinuous intervals. Moreover, the decision boundaries appear to be interleaved between successive stages, and their frequency increases as k increases. In the bottom row, we show the reconstruction function for some of the possible decoder inputs at each stage, in function of the side information. Moving towards the right-most plots highlights the successive refinement of the output for the selected bins. Similar to [16], we observe nearly linear reconstruction functions within each bin, which is the theoretically optimal strategy.

5. CONCLUSION

In this work, we designed the first learned successively refined model for Wyner-Ziv coding. We used two upper bounds on the Wyner-Ziv rate and added a distortion term to obtain two loss functions to train our RNN models. We showed that the learned successive models can attain a performance close to the state-of-the-art monolithic learned Wyner-Ziv coder, showing the viability of neural networks in successively refined coding. Additionally, we showed that the models can learn nested quantization and binning behavior, similar to the operations in the achievability proof in the Wyner-Ziv theorem.

6. REFERENCES

- [1] D. Slepian and J. Wolf, "Noiseless coding of correlated information sources," *IEEE Transactions on Information Theory*, vol. 19, no. 4, pp. 471–480, 1973.
- [2] A. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Transactions on Information Theory*, vol. 22, no. 1, pp. 1–10, 1976.
- [3] Y. Steinberg and N. Merhav, "On successive refinement for the Wyner-Ziv problem," *IEEE Transactions on Information Theory*, vol. 50, no. 8, pp. 1636–1654, 2004.
- [4] S. Pradhan and K. Ramchandran, "Distributed source coding using syndromes (DISCUS): Design and construction," *IEEE transactions on information theory*, vol. 49, no. 3, pp. 626–643, 2003.
- [5] R. Zamir, S. Shamai, and U. Erez, "Nested linear/lattice codes for structured multiterminal binning," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1250–1276, 2002.
- [6] S. Cheng and Z. Xiong, "Successive refinement for the Wyner-Ziv problem and layered code design," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 3269–3281, 2005.
- [7] N. Deligiannis, A. Munteanu, S. Wang, S. Cheng, and P. Schelkens, "Maximum likelihood Laplacian correlation channel estimation in layered Wyner-Ziv coding," *IEEE Transactions on Signal Processing*, vol. 62, no. 4, pp. 892–904, 2014.
- [8] Johannes Ballé, Valero Laparra, and Eero P Simoncelli, "End-to-end optimized image compression," in *International Conference on Learning Representations*, 2016.
- [9] G. Toderici, D. Vincent, N. Johnston, S. Jin Hwang, D. Minnen, J. Shor, and M. Covell, "Full resolution image compression with recurrent neural networks," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017, pp. 5306–5314.
- [10] R. Royen, L. Denis, Q. Bolsee, P. Hu, and A. Munteanu, "MaskLayer: Enabling scalable deep learning solutions by training embedded feature sets," *Neural Networks*, vol. 137, pp. 43–53, 2021.
- [11] J. Whang, A. Acharya, H. Kim, and A. Dimakis, "Neural distributed source coding," *arXiv preprint arXiv:2106.02797*, 2021.
- [12] N. Mital, E. Özyilkan, A. Garjani, and D. Gündüz, "Neural distributed image compression with cross-attention feature alignment," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2498–2507.
- [13] X. Zhang, J. Shao, and J. Zhang, "LDMIC: Learning-based distributed multi-view image coding," in *The Eleventh International Conference on Learning Representations*, 2022.
- [14] L. Abrahamyan, Y. Chen, G. Bekoulis, and N. Deligiannis, "Learned gradient compression for distributed deep learning," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 12, pp. 7330–7344, 2021.
- [15] E. Lei, H. Hassani, and S. Bidokhti, "Federated neural compression under heterogeneous data," in *2023 IEEE International Symposium on Information Theory (ISIT)*, 2023, pp. 525–530.
- [16] E. Özyilkan, J. Ballé, and E. Erkip, "Learned wyner-ziv compressors recover binning," in *2023 IEEE International Symposium on Information Theory (ISIT)*, 2023, pp. 701–706.
- [17] S. Pradhan, J. Chou, and K. Ramchandran, "Duality between source coding and channel coding and its extension to the side information case," *IEEE Transactions on Information Theory*, vol. 49, no. 5, pp. 1181–1203, 2003.
- [18] N. Deligiannis, A. Sechelea, A. Munteanu, and S. Cheng, "The no-rate-loss property of wyner-ziv coding in the z-channel correlation case," *IEEE Communications Letters*, vol. 18, no. 10, pp. 1675–1678, 2014.
- [19] E. Jang, S. Gu, and B. Poole, "Categorical Reparameterization with Gumbel-Softmax," in *International Conference on Learning Representations*, 2017.
- [20] C. Maddison, A. Mnih, and Y. Teh, "The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables," in *International Conference on Learning Representations*, 2017.