Self-attention for Enhanced OAMP Detection in MIMO Systems

A Preprint

Alexander Fuchs
Signal Processing and Speech Communication Lab.
Graz University of Technology

Christian Knoll
Signal Processing and Speech Communication Lab.
Graz University of Technology

Nima N. Moghadam
Huawei Technologies Sweden AB

Alexey Pak
Huawei Technologies Sweden AB

Jinliang Huang
Huawei Technologies Sweden AB

Erik Leitinger
Signal Processing and Speech Communication Lab.
Graz University of Technology

Franz Pernkopf
Signal Processing and Speech Communication Lab.
Graz University of Technology

March 15, 2023

Abstract

Multiple-Input Multiple-Output (MIMO) systems are essential for wireless communications. Since classical algorithms for symbol detection in MIMO setups require large computational resources or provide poor results, data-driven algorithms are becoming more popular. Most of the proposed algorithms, however, introduce approximations leading to degraded performance for realistic MIMO systems. In this paper, we introduce a neural-enhanced hybrid model, augmenting the analytic backbone algorithm with state-of-the-art neural network components. In particular, we introduce a self-attention model for the enhancement of the iterative Orthogonal Approximate Message Passing (OAMP)-based decoding algorithm. In our experiments, we show that the proposed model can outperform existing data-driven approaches for OAMP while having improved generalization to other SNR values at limited computational overhead.

1 Introduction

Multiple-Input Multiple-Output (MIMO) systems have become an integral part of communication standards over the past few years. They offer increased spectral and energy efficiency, as well as improved spatial resolution between users, and link reliability [1, 2]. Efficient and precise channel estimation and symbol detection is vital to leverage the benefits of MIMO systems. Although the Maximum Likelihood (ML) detector can achieve optimal performance for the detection problem, it scales exponentially with the size of the MIMO system, which limits its practical relevance. The sphere-decoding algorithm alleviates this by limiting the search space but remains intractable for large systems [3]. The Zero Forcing (ZF) detector and Linear Minimum Mean Squared Error (MMSE) detector [4], are able to solve the detection problem efficiently but do not perform as well as the ML detector.

There are two alternative directions for improving the detection quality at limited computational complexity. First, approximating the maximum likelihood solution by iterative message passing algorithms: these include Expectation Propagation (EP), Approximate Message Passing (AMP), and Orthogonal AMP (OAMP) [5, 6, 7]. Particularly the OAMP detector shows excellent performance for idealized independent and identically distributed (i.i.d.) input data and uncorrelated channels. Because of the implicit assumptions and approximations, the performance of the OAMP algorithm degrades for more realistic scenarios with correlated data and channels though. Second, spurred by the recent successes of machine learning, it is tempting to attempt learning the detection directly from data. However, given the complexity of the problem, this would require huge amounts of training data and cannot be expected to generalize

well to scenarios not captured in the training data. Instead, successful data-driven approaches in such a complex setting, need to exploit neural networks and be build upon existing iterative frameworks via algorithm unrolling [8]. The first approach of this kind was DetNet [9]; later approaches that incorporate some data dependence – such as MMNet and OAMPNet – could improve the detection performance, especially for uncorrelated channels [10, 11, 12]. All these data-driven approaches learn relatively few static parameters, and although this accelerates the convergence and improves the detection performance, it still leaves room for improvement by fully utilizing the flexibility and capabilities of neural networks. So-called neural-enhanced or hybrid models have recently been proposed with great success for Belief Propagation (BP) based detection [13]. These hybrid models rely on graph neural networks to refine the message updates within the BP iteration. The overarching goal is to combine the flexibility of neural networks with the modeling capabilities of generative models to: (i) mitigate the approximation error of the generative model, (ii) accelerate convergence, and (iii) improve the detection performance.

If we want to achieve good detection performance reliably with good generalization (i.e., also for scenarios not seen during training), we need a close integration of purely data-driven approaches with generative models that capture strong prior information about the underlying system in the form of inductive biases.

In this paper, we propose a *neural-enhanced* OAMP algorithm. More precisely, we use a self-attention model for augmenting the OAMP iteration and, in doing so, reducing the dependence on the rigid assumptions of i.i.d. data and uncorrelated channels. Self-attention was developed for transformer models [14], which are now used in a wide range of areas ranging from speech, over image, to signal-processing applications [15, 16, 17]. One of the main benefits of using a self-attention model is that it is agnostic to the number of users and permutations of the transmitted signal. This allows for a symbiotic relationship with the analytical OAMP backbone algorithm that not only improves the decoding performance but also offers excellent generalization capabilities.[1] In our experiments, we analyze the model performance and generalization capabilities on a Rayleigh channel and evaluate the symbol detection performance on more realistic 3GPP channel data.

2 Self-attention enhanced OAMP for MIMO decoding

The MIMO decoding setup for M receive and N transmit antennas is given by a linear system, i.e.,

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \boldsymbol{\nu}, \tag{1}$$

where $\mathbf{y} \in \mathbb{C}^M$ represents the received signal vector, $\mathbf{H} \in \mathbb{C}^{M \times N}$ is the channel matrix, $\mathbf{x} \in \mathbb{C}^N$ is a vector containing the transmit symbols, and $\boldsymbol{\nu} \in \mathbb{C}^M$ is a noise vector containing samples from a complex Gaussian distribution $\mathcal{CN}(0, \sigma)$ with standard deviation σ. In this work, we assume that $\mathbf{H}$ and σ are known or estimated by a preceding channel estimation step. The decoder has to determine an inverse solution of Eq. 1 by detecting the correct transmitted symbols $\mathbf{x}$. The optimal estimate $\hat{\mathbf{x}}$ of the transmit symbols $\mathbf{x}$ is given by the maximum likelihood estimator, i.e.,

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} \mathcal{P}(\mathbf{x}|\mathbf{y}, \mathbf{H}) \tag{2}$$

with

$$\mathcal{P}(\mathbf{x}|\mathbf{y}, \mathbf{H}) = \frac{\mathcal{P}(\mathbf{y}|\mathbf{x}, \mathbf{H})\mathcal{P}(\mathbf{x})}{\mathcal{P}(\mathbf{y}|\mathbf{H})}. \tag{3}$$

Here, the data evidence $\mathcal{P}(\mathbf{y}|\mathbf{H})$ is obtained by marginalization over $\mathbf{x}$,

$$\mathcal{P}(\mathbf{y}|\mathbf{H}) = \sum_{\mathbf{x} \in \mathcal{S}} \mathcal{P}(\mathbf{y}|\mathbf{x}, \mathbf{H})\mathcal{P}(\mathbf{x}). \tag{4}$$

However, the estimate $\hat{\mathbf{x}}$ is in general computationally intractable as it involves the evaluation of high-dimensional integrals (see Eq. 2 and Eq. 4) [4]. Therefore, several algorithms have been proposed to improve computational efficiency by approximating $\mathcal{P}(\mathbf{x}|\mathbf{y}, \mathbf{H})$ [6, 7, 18]. OAMP-based algorithms have shown to perform well for symbol detection. They decouple the posterior pdf by assuming independent Gaussian distributed symbols $\mathbf{x}$. This leads to an iterative state evolution algorithm that, if converged, has the potential to retrieve the optimal symbol estimates $\hat{\mathbf{x}}$ for independent and identically distributed data and uncorrelated channels. The OAMP algorithm is defined in Algorithm 1.

The OAMPNet models are both based on the OAMP algorithm [11, 12], but include additional scaling parameters, which improve the convergence of the algorithm, leading to better results at the same number of iterations. The OAMPNet2 model uses the trainable parameter γ to scale the updates of the linear estimate $\mathbf{r}_t$, exchanging Eq.7 for:

$$\mathbf{r}_t \leftarrow \mathbf{x}_{t-1} + \gamma_t \mathbf{W}_t(\mathbf{y} - \mathbf{H}\mathbf{x}_{t-1}) \tag{12}$$

---

[1]Code: https://github.com/alexf1991/self_attention_oamp_mimo

---

Algorithm 1 OAMP algorithm for MIMO detection

    Input: Channel matrix $\mathbf{H}$, received signal $\mathbf{y}$ and noise
    standard deviation σ, number of iterations T,
    number of receive antennas M, number of symbols N.
    Output: Estimated symbols $\mathbf{x}_T$.
    Initialize: $\tau_{t-1} \leftarrow 1$, $\mathbf{x}_{t-1} \leftarrow 0$, $v_{t-1} \leftarrow 1$
  for $t < T$ **do**

$$\hat{\mathbf{W}}_t \leftarrow v_{t-1}^2 \mathbf{H}^T \left(v_{t-1}^2 \mathbf{H}\mathbf{H}^T + \frac{\sigma^2}{2}\mathbf{I} \right)^{-1} \tag{5}$$

$$\mathbf{W}_t \leftarrow \frac{2N}{\mathrm{tr}(\hat{\mathbf{W}}_t \mathbf{H})} \tag{6}$$

$$\mathbf{r}_t \leftarrow \mathbf{x}_{t-1} + \mathbf{W}_t(\mathbf{y} - \mathbf{H}\mathbf{x}_{t-1}) \tag{7}$$

$$\mathbf{x}_t \leftarrow \mathbb{E}\{\mathbf{s}|\mathbf{r}_t, \tau_{t-1}\} \tag{8}$$

$$v_t^2 \leftarrow \frac{||\mathbf{y} - \mathbf{H}\mathbf{x}_t||_2^2 - M\sigma^2}{\mathrm{tr}(\mathbf{H}^T\mathbf{H})} \tag{9}$$

$$\mathbf{B} \leftarrow \mathbf{I} - \mathbf{W}_t\mathbf{H} \tag{10}$$

$$\tau_t^2 \leftarrow \frac{1}{2N}\mathrm{tr}(\mathbf{B}\mathbf{B}^T)v_t^2 + \frac{1}{4N}\mathrm{tr}(\mathbf{W}_t\mathbf{W}_t^T)\sigma^2 \tag{11}$$

  end for

---

and a second parameter θ as a scaling parameter for $\mathbf{W}$, resulting in a modified $\mathbf{B}$ in Eq. 10, and variance τ^2 in Eq. 11 for the non-linear estimator:

$$\mathbf{B} \leftarrow \mathbf{I} - \theta_t\mathbf{W}_t\mathbf{H},$$
$$\tau_t^2 \leftarrow \frac{1}{2N}\mathrm{tr}(\mathbf{B}\mathbf{B}^T)v_t^2 + \frac{\theta^2}{4N}\mathrm{tr}(\mathbf{W}_t\mathbf{W}_t^T)\sigma^2. \tag{13}$$

Furthermore, it includes the linear estimate $\mathbf{r}_t$ in addition to the non-linear estimate $\mathbb{E}\{\mathbf{s}|\mathbf{r}_t, \tau_t\}$ to create $\mathbf{x}_{t+1}$, using two additional trainable parameters ϕ_t and ζ_t as weighting factors, and the symbol alphabet $\mathcal{S}$ given by the modulation scheme. The updated symbol estimates are therefore given by

$$\mathbf{x}_t = \phi_t(\mathbb{E}\{\mathbf{s}|\mathbf{r}_t, \tau_{t-1}\} - \zeta_t\mathbf{r}_t). \tag{14}$$

OAMPNet2 performs well for i.i.d data and weakly correlated channels, however, the detection performance degrades significantly in more realistic scenarios characterized by stronger correlations between the individual MIMO channels as the approximations introduce a model mismatch. The goal of our proposed neural-enhanced hybrid model is to (i) reduce the impact of correlated channels on the detection performance by mitigating the model mismatch and (ii) to improve the convergence behavior of the algorithm. Both measures lead to improved detection performance. Our proposed model is based on the OAMP algorithm and extends the algorithm with a data-driven self-attention model to augment the predictions.

2.1 Self-attention

The concept of self-attention was first introduced for transformer models in natural language processing [14]. Transformers are a special type of sequence-to-sequence DNNs, which were introduced to improve the contextual understanding of words. They implement a multi-stage encoder-decoder structure, where each stage includes a multi-headed attention layer, that uses self-attention to establish relation between words. Here, each word is represented by a vector token $\mathbf{z}_i$, where $i \in \{1, 2 \ldots T\}$ and T is the number of tokens.

The self-attention layer creates multiple projections of the tokens that are then used to create an attention matrix $\mathbf{A}$ which intends to represent the relations between tokens. This attention matrix is based on scalar products between the tokens and can adjust according to the input sentence. To create these projections, three different linear layers are applied to every input token $\mathbf{z}_i$. Since each of the three resulting vectors fulfills a different role within the self-attention layer, they are referred to as:

- *key* vector: $\mathbf{k}_i = \mathbf{K}\mathbf{z}_i$,
- *query* vector: $\mathbf{q}_i = \mathbf{Q}\mathbf{z}_i$,
- *value* vector: $\mathbf{v}_i = \mathbf{V}\mathbf{z}_i$,

where $\mathbf{K}$, $\mathbf{Q}$, and $\mathbf{V}$ are linear projection matrices, that are learned during training. This naming convention originates from the sequence-to-sequence literature, as there the decoder queries a key for a specific token from the encoder, or vice versa, to perform matching between tokens. The intention of the *key* and *query* vectors are to form connections between tokens that can benefit from information exchange and to decouple tokens that are independent. The elements of the attention matrix $A_{i,j}$ are then created by first taking the absolute value of the scalar product between the key and query vectors of each input $A_{i,j} = |\mathbf{k}_i^{\mathrm{T}} \mathbf{q}_j|$, before a $\mathrm{softmax}$ function, using an inverse temperature parameter β, is applied:

$$\bar{A}_{i,j} = \operatorname*{softmax}_{j} \left(\beta A_{ij} \right). \tag{15}$$

Here, the inverse temperature is typically chosen as one over the square root of the number of used vector dimensions $\beta = 1/\sqrt{d_k}$, where $d_k = \dim(\mathbf{k})$. The normalized attention matrix $\bar{\mathbf{A}}$ is then applied to the value vector tokens $\mathbf{v}_j$ to create the output tokens via a weighted sum,

$$\bar{\mathbf{v}}_i = \sum_j \bar{A}_{i,j} \mathbf{v}_j. \tag{16}$$

Afterwards two layers $\mathbf{W}_1 \in \mathbb{R}^{d_k \times d_k}$, $\mathbf{W}_2 \mathbb{R}^{d_k \times d_k}$, layer normalization and a ReLU activation function are applied to create the final output vector tokens $\bar{\mathbf{z}}_i$ of the self-attention [14]. Layer normalization (LN) for the latent vector $\mathbf{z}$ is defined as

$$\mathrm{LN}(\mathbf{z}) = \frac{\mathbf{z} - \hat{\mu}}{\hat{\sigma}}, \hat{\mu} = \frac{1}{C_l} \sum_c^{C_l} z_i, \hat{\sigma} = \sqrt{\frac{1}{C_l} \sum_c^{C_l} (z_i - \hat{\mu})^2}, \tag{17}$$

where C_l is the total number of channels, i.e. vector dimensions c. Each layer produces an output of the same dimension as the previous layer to enable the use of shortcut connections and prevent vanishing gradients:

$$\begin{aligned} \mathbf{a}_i &= \mathrm{LN}(\mathbf{z}_i + \bar{\mathbf{v}}_i), \\ \mathbf{b}_i &= \mathbf{W}_2 \mathrm{ReLU}(\mathbf{W}_1 \mathbf{a}_i), \\ \bar{\mathbf{z}}_i &= \mathrm{LN}(\mathbf{b}_i + \mathbf{a}_i). \end{aligned} \tag{18}$$

2.2 Self-attention for MIMO detection

For MIMO detection algorithms independence from the user order, and the total number of users is essential. Self-attention models can explicitly include these properties in the model if no position encodings are used. This helps the model to learn more efficiently and improve generalization, similar to the translation invariant kernel functions in Convolutional Neural Networks (CNNs). In our proposed setup, the self-attention model acts as a post-processing step for the linear estimator of the OAMP iteration. Meaning, that the OAMP algorithm first generates estimates for $\mathbf{r}_t$ and v_t^2, which are then included in the input for the self-attention model. Here, we omit Eq. 8, Eq. 10, and Eq. 11. Instead, the model predicts symbol likelihood estimates $\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t$ for each of the K symbols s_k in the alphabet defined via the modulation scheme, which are then used to calculate the expected symbol estimates.:

$$\mathbf{x}_t \leftarrow \mathbb{E}\{\mathbf{x} \in \mathcal{S}|\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t\}. \tag{19}$$

Additionally, the model also predicts a refined the variance estimate $\hat{v}_t^2 = g(f(\mathbf{z}_{i,t}))$. For appropriate data representation, we create two input vector tokens per user, representing real- and imaginary parts. We compose symbol and channel information for each user in a vector, and combine it with the current OAMP estimates for $\mathbf{r}_t$, $\boldsymbol{v}_t^2$, and the symbol likelihood estimates of the previous iteration $\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_{t-1}$.

2.2.1 Data preparation

In order to optimize the model for a MIMO detection task we structure the input data into $2N$ vector tokens, with two tokens for the real- and imaginary part of each user. This data structure preserves all information while staying limited in size and computational complexity, even for large MIMO systems. Each input vector token $\mathbf{z}_i$ consists of the following components:

- The current latent state of the transformer network in the form of a latent vector $\bar{\mathbf{b}}_{i,t-1} \in \mathbb{R}^{d_b}$, where $d_{\bar{b}} = dim(\bar{\mathbf{b}}_{i,t-1})$ is the number of used state channels.
- The received signal $\mathbf{y} \in \mathbb{R}^M$.
- The row of channel matrix corresponding to user n $\mathbf{H}_n \in \mathbb{R}^M$.
- The last estimate for $\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_{t-1} \in \mathbb{R}^K$.
- The last received symbol estimate $x_{t-1} \in \mathbb{R}^1$.
- The current estimate for $r_t \in \mathbb{R}^1$.
- The current estimate for $\hat{v}_t^2 \in \mathbb{R}^1$.

The final input vectors $\mathbf{z}_{i,t}$ are created by concatenating the vector components above along the channel dimension and normalizing using layer normalization. This results in a vector $\mathbf{z}_i \in R^C$, with $C = 2M + K + d_{\bar{b}} + 3$.

2.2.2 Estimation

The self-attention model $f(\mathbf{z}_{i,t})$ then processes the tokens into latent vectors $\mathbf{b}_{i,t} = f(\mathbf{z}_{i,t}) \in \mathbb{R}^C$, which are then used for the estimation of $\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t$ and $\hat{v}_t^2$. The first $d_{\hat{b}}$ dimensions of $\mathbf{b}_{i,t}$ are also used as the new latent state $\bar{\mathbf{b}}_{i,t}$. For the likelihood estimation of the alphabet, two dense layers are applied to each $\mathbf{b}_{i,t}$:

$$\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t = \sigma(\mathbf{W}_2 \text{LN}(\text{ReLU}(\mathbf{W}_1 \mathbf{b}_{i,t}))), \tag{20}$$

where $\mathbf{W}_1 \in \mathbb{R}^{C \times d_{\bar{b}}}$, $\mathbf{W}_2 \in \mathbb{R}^{d_{\bar{b}} \times 1}$ are dense layers, LN is layer normalization, and $\sigma(\cdot)$ is the logistic sigmoid activation function. For the estimation of $\hat{v}_t^2$, we first average $\mathbf{b}_i$ over the tokens, and then apply a single dense layer $\mathbf{W}_3 \in \mathbb{R}^{d_{\bar{b}} \times 1}$ reducing the vector to a scalar. Then a squared logistic sigmoid activation function is applied to create $\hat{v}_t^2$:

$$\hat{v}_t^2 = \sigma\left(\mathbf{W}_3\left(\frac{1}{2N}\sum_i^{2N}\mathbf{b}_{i,t}\right)\right)^2. \tag{21}$$

3 Experiments

In our experiments, we evaluate the performance and generalization behavior as well as the runtime of our proposed OAMP-Self-Attention (OAMP-SA) algorithm and compare it against OAMP, OAMPNet2, and the MMSE estimator without interference cancellation according to [4]. Section 3.1 provides a comparison of the generalization behavior of the different algorithms; in Section 3.2 we analyze the performance in a more realistic MIMO scenario using 3GPP channel matrices; and Section 3.3 analyzes the runtime for large MIMO systems. The data-driven models OAMP-SA and OAMPNet2 are trained using a cross-entropy loss for each iteration:

$$\mathcal{L}_{CE} = \sum_t \sum_k -(y_k \log(\mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t) + (1 - y_k) \log(1 - \mathcal{P}(s_k|f(\mathbf{z}_{i,t}))_t)), \tag{22}$$

where y_k is a binary label for the specific symbol probability. We perform the training of the model for 25 epochs using an ADAM optimizer with a learning rate of $lr = 0.001$ [19]. After 10 and 20 epochs we decay the learning rate by a factor of 10. All iterative algorithms are evaluated after 10 iterations.

3.1 Model Generalization

In this experiment, we investigated a MIMO system using $N_t = 8$ transmit antennas and $N_r = 16$ receive antennas for a Gaussian i.i.d channel with Rayleigh fading using a QAM-4 encoding. The datasets consist of 80k channel realizations for training, as well as a validation set, and test set with 20k channel realization, for each integer value within the SNR range of 5-14 dB. To analyze generalization for the data-driven models, we evaluate two settings: (i) Training the models on MIMO data with a specific SNR and specific correlation coefficients c_r, c_t, and then evaluating on data with the same SNR and correlation, (i.e., the non-generalizing setting) and (ii) training the model on data with a fixed SNR value of 10 dB an with fixed correlation coefficients $c_r = c_t = 0.1$, and then evaluating the same model over the entire SNR range and at multiple correlations (i.e., the generalizing setting).

In Figure 1, we see the Symbol Error Rate (SER) of the model at $c_t = c_r = 0.1$, varying the SNR of the MIMO data. Here, the self-attention model manages to substantially outperform OAMPNet2 over the entire SNR range. Interestingly, we also see that for the high SNR region, the OAMP-SA model in setting (ii) outperforms variant (i). This underlines the excellent generalization capabilities of the model: although our proposed model is more expressive than the alternative hybrid models, it generalizes well when evaluated on scenarios with SNR values not covered in the training data. This also indicates that for data-driven models training data selection is crucial, and sufficiently difficult scenarios need to be covered during training.

3.2 Realistic channel data

For this evaluation, we use a more realistic scenario, utilizing 3GPP channel realizations generated via a clustered delay line (CDL)-B model [20]. The investigated MIMO system uses $N_r = 64$ antennas, with beam selection. After beam selection, the resulting channel matrices contain $\bar{N}_r = 16$ beams, for a channel recorded over 200 time steps using 96 subcarriers. The selection procedure randomly chooses $N_t = 12$ users out of a total of 100, creating 20 sets of users for training and another 20 sets for the test set. From the training set, we split 20 % as a validation set, resulting
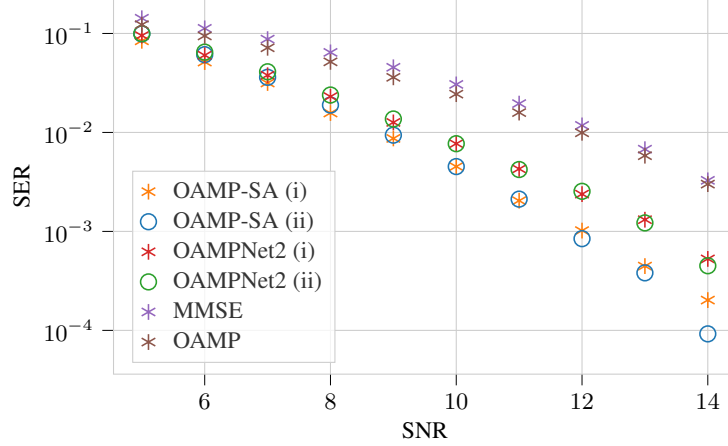
Figure 1: Evaluation of different decoding algorithms over SNR for a 16×8 QAM 4 MIMO system using the SER metric. Stars mark evaluations for setting (i) (non-generalizing) with $c_r = c_t = 0.1$; circles mark evaluations for setting (ii) (generalizing) with varying data SNR.

in $16 \times 200 \times 96$ channel matrices for training, $4 \times 200 \times 96$ for validation, and $20 \times 200 \times 96$ matrices for testing. Here, the training and test set use two disjoint sets of users in order to increase the diversity between the training and test set in order to detect overfitting on specific scenarios. Each channel is then evaluated for two symbol transmissions, for training as well as for testing. The transmissions for the training procedure were randomly drawn from an SNR range of 5-20 dB. Figure 2 shows, that our proposed OAMP-SA model can also improve performance on realistic channel
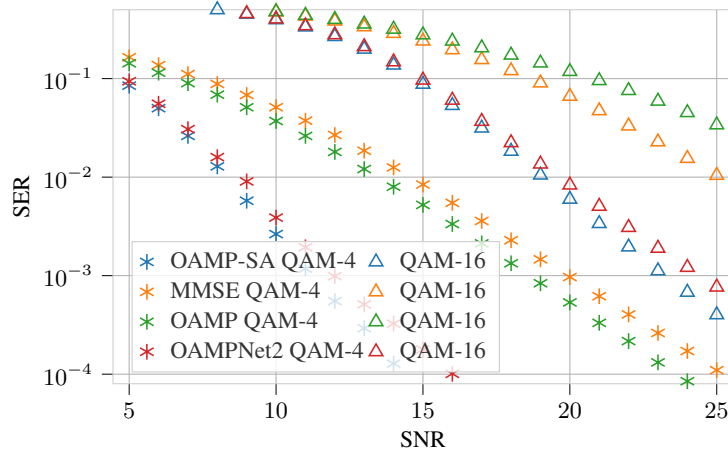

Figure 2: Algorithm comparison on a 3GPP channel scenario after beam selection. The system uses $\bar{N}_r = 16$ beams, for $N_t = 12$ users and QAM-4 (stars) as well as QAM-16 (triangles) modulation.

data, and does not degrade in performance for the SNR region not covered in the training distribution ranging from 20-25 dB SNR. Interestingly, OAMP performance for the QAM-16 system degrades compared to MMSE for the high SNR region, which as a result might also affect OAMPNet2 and OAMP-SA performance.

3.3 Runtime evaluation

In this experiment, we evaluate the average runtime of the MMSE, OAMP, OAMPNet2, and OAMP-SA algorithms on a GPU system using an NVIDIA RTX 2080Ti. Here, we investigate two systems of size $N_r = 64$, $N_t = 32$, and $N_r = 512$, $N_t = 256$. These system sizes allow for a better assessment of the algorithm's scalability.

In Table 1, we see that the OAMP-SA model takes three times as long as OAMPNet2 per channel evaluation for the smaller 64×32 system, and about two times as long for the larger 512×256 system. These results show, that while the self-attention model does add some complexity, it scales well to larger systems.

Table 1: Average runtime of the investigated algorithms for a $N_r = 64$, $N_t = 32$ and a $N_r = 512$, $N_t = 256$ QAM-4 MIMO system on GPU hardware.

System	MMSE	OAMP	OAMPNet2	OAMP-SA
64×32	1.6 μs	10.9 μs	11.9μs	32.8μs
512×256	47.0μs	93.8μs	125.0μs	237.5μs

4 Conclusion

In this paper, we proposed a self-attention model to enhance the OAMP algorithm for a realistic MIMO system. The introduced deep learning architecture is able to incorporate an analytical prior that makes it more data efficient and improves generalization at limited computational overhead. In our experiments, we showed that our model is able to generalize well across a wide SNR range, and is able to outperform existing OAMP-based approaches on realistic channel data while scaling well to large MIMO systems. For future investigations, we would like to explore the applicability of self-attention models for other iterative decoding algorithms.

References

[1] Andrea Goldsmith, Syed Ali Jafar, Nihar Jindal, and Sriram Vishwanath, "Capacity limits of MIMO channels," *IEEE Journal on selected areas in Communications*, vol. 21, no. 5, pp. 684–702, 2003.

[2] Erik G Larsson, Ove Edfors, Fredrik Tufvesson, and Thomas L Marzetta, "Massive MIMO for next generation wireless systems," *IEEE communications magazine*, vol. 52, no. 2, pp. 186–195, 2014.

[3] Zhan Guo and Peter Nilsson, "Algorithm and implementation of the k-best sphere decoding for MIMO detection," *IEEE Journal on selected areas in communications*, vol. 24, no. 3, pp. 491–503, 2006.

[4] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and challenges with very large arrays," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[5] Javier Cespedes, Pablo M Olmos, Matilde Sánchez-Fernández, and Fernando Perez-Cruz, "Expectation propagation detection for high-order high-dimensional MIMO systems," *IEEE Transactions on Communications*, vol. 62, no. 8, pp. 2840–2849, 2014.

[6] David L Donoho, Arian Maleki, and Andrea Montanari, "Message-passing algorithms for compressed sensing," *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18914–18919, 2009.

[7] Junjie Ma and Li Ping, "Orthogonal amp," *IEEE Access*, vol. 5, pp. 2020–2033, 2017.

[8] Vishal Monga, Yuelong Li, and Yonina C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.

[9] Neev Samuel, Tzvi Diskin, and Ami Wiesel, "Deep MIMO detection," in *2017 IEEE 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 2017, pp. 1–5.

[10] Mehrdad Khani, Mohammad Alizadeh, Jakob Hoydis, and Phil Fleming, "Adaptive neural signal detection for massive MIMO," *IEEE Transactions on Wireless Communications*, vol. 19, no. 8, pp. 5635–5648, 2020.

[11] Hengtao He, Chao-Kai Wen, Shi Jin, and Geoffrey Ye Li, "A model-driven deep learning network for MIMO detection," in *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2018, pp. 584–588.

[12] Hengtao He, Chao-Kai Wen, Shi Jin, and Geoffrey Ye Li, "Model-driven deep learning for MIMO detection," *IEEE Transactions on Signal Processing*, vol. 68, pp. 1702–1715, 2020.

[13] Victor Garcia Satorras and Max Welling, "Neural enhanced belief propagation on factor graphs," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 685–693.

[14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.

[15] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12449–12460, 2020.

[16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al., "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.

[17] Juliano Pinto, Georg Hess, William Ljungbergh, Yuxuan Xia, Henk Wymeersch, and Lennart Svensson, "Can deep learning be applied to model-based multi-object tracking?," *arXiv preprint arXiv:2202.07909*, 2022.

[18] Gunvor Elisabeth Kirkelund, Carles Navarro Manchón, Lars PB Christensen, Erwin Riegler, and Bernard Henri Fleury, "Variational message-passing for joint channel estimation and decoding in MIMO-ofdm," in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*. IEEE, 2010, pp. 1–6.

[19] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[20] TR ETSI, "138 901 v14. 3.0: 5g; study on channel model for frequencies from 0.5 to 100 ghz, 3gpp std," Tech. Rep., 3GPP TR 38.901 version 14.0. 0 release 14, 2017.