

PREDICTING FLAT-FADING CHANNELS VIA META-LEARNED CLOSED-FORM LINEAR FILTERS AND EQUILIBRIUM PROPAGATION

Sangwoo Park Osvaldo Simeone

King’s Centre for Learning & Information Processing (KCLIP),
Department of Engineering, King’s College London

ABSTRACT

Predicting fading channels is a classical problem with a vast array of applications, including as an enabler of artificial intelligence (AI)-based proactive resource allocation for cellular networks. Under the assumption that the fading channel follows a stationary complex Gaussian process, as for Rayleigh and Rician fading models, the optimal predictor is linear, and it can be directly computed from the Doppler spectrum via standard linear minimum mean squared error (LMMSE) estimation. However, in practice, the Doppler spectrum is unknown, and the predictor has only access to a limited time series of estimated channels. This paper proposes to leverage meta-learning in order to mitigate the requirements in terms of training data for channel fading prediction. Specifically, it first develops an offline low-complexity solution based on linear filtering via a meta-trained quadratic regularization. Then, an online method is proposed based on gradient descent and equilibrium propagation (EP). Numerical results demonstrate the advantages of the proposed approach, showing its capacity to approach the genie-aided LMMSE solution with a small number of training data points.

Index Terms— Meta-learning, ridge regression, fading channels, equilibrium propagation

1. INTRODUCTION

The capacity to accurately predict channel state information (CSI) is an enabler of efficient resource allocation strategies. Notably, it allows the transmitter of a communication link to optimize modulation and coding schemes, while reducing the pilot overhead; and it can be leveraged by artificial-intelligence (AI) tools to implement proactive networking protocols [1]. Given that standard fading channel models follow stationary complex Gaussian processes, e.g., for Rayleigh and Rician fading, optimal channel predictors can be obtained via linear minimum mean squared error (LMMSE) estimators, such as the Wiener filter [2]. This, however, requires the channel statistics, including the Doppler spectrum, to be known. In practice, CSI statistics are a priori unknown, and need to be estimated from data.

The work of O. Simeone was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 725731).

Accordingly, there is a long line of work that focuses on predicting time-varying channels with linear predictors obtained via autoregressive (AR) models [3] or Kalman filtering [4] by estimating channel statistics from available pilot data. Recently, deep learning-based nonlinear predictors have been proposed to adapt to channel statistics through the training of neural networks [5, 6]. However, as reported in [5, 6], deep learning based predictors tend to require excessive training (pilot) data, while failing to outperform well-designed linear filters in the low-data regime.

In this paper, we propose to leverage meta-learning in order to further reduce the data requirement of linear CSI predictors for flat-fading channels. The key idea is to leverage multiple related channel data sets from various propagation environments to optimize the hyperparameters of a training procedure. Specifically, we first develop a low-complexity offline solution that obtains hyperparameters and predictors in closed form. Then, an online algorithm is proposed based on gradient descent and equilibrium propagation (EP) [7, 8]. Previous applications of meta-learning to communication systems include demodulation [9, 10], channel equalization [11], encoding/decoding [12, 13, 14], MIMO detection [15], beamforming [16, 17], and resource allocation [18].

2. SYSTEM MODEL AND BACKGROUND

Notation: In this paper, $(\cdot)^\top$ denotes transposition; $(\cdot)^\dagger$ Hermitian transposition; $|\cdot|$ the absolute value; $\|\cdot\|$ the Euclidean norm; and I the identity matrix.

2.1. Preliminaries

In order to establish some useful notation, consider a supervised learning problem targeting the prediction of a scalar label y given a $d \times 1$ covariate vector x . Both x and y are complex-valued. The training data set $\mathcal{Z} = (X, y)$ includes n labelled examples $\{(x_i, y_i)\}_{i=1}^n$. Specifically, the $n \times d$ matrix $X = [x_1^\dagger, \dots, x_n^\dagger]^\top$ contains by row the Hermitian transpose of the $d \times 1$ covariate vectors $\{x_i\}_{i=1}^n$, while the $n \times 1$ vector $y = [y_1, \dots, y_n]^\dagger$ collects the complex conjugate of the corresponding labels $\{y_i\}_{i=1}^n$. A linear predictor $\hat{y} = v^\dagger x$ of the label y is defined by a $d \times 1$ model parameter vector v .

Given data set \mathcal{Z} and hyperparameters (λ, \bar{v}) , with $\lambda > 0$ and $d \times 1$ bias vector \bar{v} , we define the solution of ridge regression as




Fig. 1: Illustration of the frame-based transmission system under study: At any frame f , based on the previous N channels $h_{l,f}^N$, we investigate the problem of optimizing the δ -lag prediction $\hat{h}_{l+\delta,f}$.

$$\begin{aligned} v_\lambda^*(\mathcal{Z}|\bar{v}) &= \arg \min_v \left\{ \sum_{i=1}^n |v^\dagger x_i - y_i|^2 + \lambda \|v - \bar{v}\|^2 \right\} \\ &= \arg \min_v \left\{ \|Xv - y\|^2 + \lambda \|v - \bar{v}\|^2 \right\}, \end{aligned} \quad (1)$$

which can be obtained explicitly as

$$v_\lambda^*(\mathcal{Z}|\bar{v}) = A^{-1}(X^\dagger y + \lambda \bar{v}), \text{ with } A = X^\dagger X + \lambda I. \quad (2)$$

2.2. System Model and Problem Definition

As illustrated in Fig. 1, we study a frame-based transmission system, with each frame containing multiple time slots. In each frame $f = 1, 2, \dots$, the channels $h_{l,f} \in \mathcal{C}$ follow a stationary process across the time slots $l = 1, 2, \dots$, with a generally distinct Doppler spectrum for each frame f . For instance, in a frame f , we may have a slow-moving user subject to time-invariant fading, while in another the channel may have fast temporal variations described by Clarke's model with a large Doppler frequency.

For each frame f , we are interested in addressing the lag- δ channel prediction problem, in which channel $h_{l+\delta,f}$ is predicted based on the N past channels $h_{l,f}^N = [h_{l,f}, \dots, h_{l-N+1,f}]^T \in \mathcal{C}^{N \times 1}$ (see Fig. 1). To this end, we adopt linear prediction with regressor $v_f \in \mathcal{C}^{N \times 1}$, so that the prediction is given as

$$\hat{h}_{l+\delta,f} = v_f^\dagger h_{l,f}^N. \quad (3)$$

The focus on linear prediction is justified by the optimality of linear estimation for Gaussian stationary processes [19], which provide standard models for fading channels in rich scattering environments [20].

To train the linear regressor, for any frame, we have available the training set $\mathcal{Z}_f^{\text{tr}} = \{(x_{i,f}, y_{i,f})\}_{i=1}^{L^{\text{tr}}} \equiv \{(h_{l,f}^N, h_{l+\delta,f})\}_{l=N}^{L^{\text{tr}}+N-1}$ encompassing L^{tr} input-output examples for the predictor (2). Note that data set $\mathcal{Z}_f^{\text{tr}}$ can be constructed from $L^{\text{tr}} + N + \delta - 1$ channels $\{h_{1,f}, \dots, h_{L^{\text{tr}}+N+\delta-1,f}\}$. In practice, these channel vectors $h_{l,f}$ are estimated using pilot symbols, and estimation noise can be easily incorporated in the model. We define the corresponding $L^{\text{tr}} \times N$ input matrix $X_f^{\text{tr}} = [x_{1,f}^\dagger, \dots, x_{L^{\text{tr}},f}^\dagger]^\top$ and $L^{\text{tr}} \times 1$ target vector $y_f^{\text{tr}} = [y_{1,f}, \dots, y_{L^{\text{tr}},f}]^\dagger$, so that the data set can be expressed as the pair $\mathcal{Z}_f^{\text{tr}} = (X_f^{\text{tr}}, y_f^{\text{tr}})$.

Based on the training data $\mathcal{Z}_f^{\text{tr}}$, for fixed hyperparameters (λ, \bar{v}) , an optimized predictor (2) is obtained via the ridge regression problem (1) with X_f^{tr} and the y_f^{tr} in lieu of X and y . This yields the regressor $v_\lambda^*(\mathcal{Z}_f^{\text{tr}}|\bar{v})$.

3. FAST TRAINING VIA OFFLINE META-LEARNING

An important problem with the linear predictor described in Sec. 2.2 is that it can only rely on data from the same frame, which may be limited due to the spectral efficiency reduction caused by the transmission of pilot symbols. In this section, we propose a solution to this problem by leveraging data from different frames via offline meta-learning, where such data is available at once in a batch manner. The next section will explore an online solution.

3.1. Meta-Training

Meta-learning operates in two distinct phases: meta-training and meta-testing. During meta-training, we have access to channel data for F frames, with $L + N + \delta - 1$ channels $\{h_{1,f}, \dots, h_{L+N+\delta-1,f}\}$ available for each frame, with $L > L^{\text{tr}}$. Following the standard meta-learning approach (e.g., [21]), we split the available L data points into L^{tr} training pairs $\{(x_{i,f}, y_{i,f})\}_{i=1}^{L^{\text{tr}}} \equiv \{(x_{i,f}^{\text{tr}}, y_{i,f}^{\text{tr}})\}_{i=1}^{L^{\text{tr}}} = \mathcal{Z}_f^{\text{tr}}$ and $L^{\text{te}} = L - L^{\text{tr}}$ test pairs $\{(x_{i,f}, y_{i,f})\}_{i=L^{\text{tr}}+1}^L \equiv \{(x_{i,f}^{\text{te}}, y_{i,f}^{\text{te}})\}_{i=1}^{L^{\text{te}}} = \mathcal{Z}_f^{\text{te}}$, resulting in two separate data sets $\mathcal{Z}_f^{\text{tr}}$ and $\mathcal{Z}_f^{\text{te}}$. The covariate vectors $\{x_{i,f}\}$ and labels $\{y_{i,f}\}$ are defined as described in Sec. 2.2 as a function of the prediction window N and prediction lag δ . We correspondingly define the $L^{\text{tr}} \times N$ input matrix X_f^{tr} and the $L^{\text{tr}} \times 1$ target vector y_f^{tr} , as well as the $L^{\text{te}} \times N$ input matrix X_f^{te} and the $L^{\text{te}} \times 1$ target vector y_f^{te} .

We propose to optimize the hyperparameter vector \bar{v} , for a fixed $\lambda > 0$, based on meta-training data. This is done so as to minimize the sum-loss of the predictors defined by the trained regressors $v_\lambda^*(\mathcal{Z}_f^{\text{tr}}|\bar{v})$ across all frames $f = 1, \dots, F$. Accordingly, estimating the loss in each frame f via the test set $\mathcal{Z}_f^{\text{te}}$ yields the meta-training problem

$$\bar{v}_\lambda^* = \arg \min_{\bar{v}} \left\{ \sum_{f=1}^F \sum_{i=1}^{L^{\text{te}}} |v_\lambda^*(\mathcal{Z}_f^{\text{tr}}|\bar{v})^\dagger x_{i,f}^{\text{te}} - y_{i,f}^{\text{te}}|^2 \right\}. \quad (4)$$

As studied in [22], the minimization in (4) is a least squares problem that can be solved in closed form as

$$\bar{v}_\lambda^* = \arg \min_{\bar{v}} \sum_{f=1}^F \left\| \tilde{X}_f^{\text{te}} \bar{v} - \tilde{y}_f^{\text{te}} \right\|^2 = (\tilde{X}^\dagger \tilde{X})^{-1} \tilde{X}^\dagger \tilde{y}, \quad (5)$$

where $L^{\text{te}} \times N$ matrix \tilde{X}_f^{te} contains by row the Hermitian transpose of the $N \times 1$ pre-conditioned input vectors $\{\lambda(A_f^{\text{tr}})^{-1} x_{i,f}^{\text{te}}\}_{i=1}^{L^{\text{te}}}$, with $A_f^{\text{tr}} = (X_f^{\text{tr}})^\dagger X_f^{\text{tr}} + \lambda I$; \tilde{y}_f^{te} is $L^{\text{te}} \times 1$ vector containing vertically the complex conjugate of the transformed outputs $\{(y_{i,f}^{\text{te}} - (y_f^{\text{tr}})^\dagger X_f^{\text{tr}} (A_f^{\text{tr}})^{-1} x_{i,f}^{\text{te}})\}_{i=1}^{L^{\text{te}}}$; the $FL^{\text{te}} \times N$ matrix $\tilde{X} = [\tilde{X}_1^{\text{te}}, \dots, \tilde{X}_F^{\text{te}}]^\top$ stacks vertically the $L^{\text{te}} \times N$ matrices $\{\tilde{X}_f^{\text{te}}\}_{f=1}^F$; and the $FL^{\text{te}} \times 1$ vector $\tilde{y} = [\tilde{y}_1^{\text{te}}, \dots, \tilde{y}_F^{\text{te}}]^\top$ stacks vertically the $L^{\text{te}} \times 1$ vectors $\{\tilde{y}_f^{\text{te}}\}_{f=1}^F$.

Unlike standard meta-learning algorithms used by most papers on communications [9, 10, 11, 12, 13, 16], such as model-agnostic meta-learning (MAML), the proposed meta-learning procedure is not based on iterative processing such as gradient descent, and it does not require the tuning of hyperparameters such as a learning rate schedule. We will discuss below how to set the only free hyperparameter λ .

Reference [22] utilizes an additional validation dataset to optimize the scalar hyperparameter λ via a cross-validation approach [23]. Given that this would require the transmission of additional frames or prevent the full use of the available frames for meta-learning, we propose to tune λ by addressing the problem

$$\lambda^* = \arg \min_{\lambda \in \Lambda} \left\{ \sum_{f=1}^F \sum_{i=1}^{L^{\text{te}}} |v_{\lambda}^*(\mathcal{Z}_f^{\text{tr}} | \bar{v}_{\lambda}^*)^{\dagger} x_{i,f}^{\text{te}} - y_{i,f}^{\text{te}}|^2 \right\}, \quad (6)$$

over a predefined discrete search Λ space for λ . In order to mitigate overfitting that may result from optimizing both hyperparameters \bar{v} and λ based on the same split $(\mathcal{Z}_f^{\text{tr}}, \mathcal{Z}_f^{\text{te}})$ of data set \mathcal{Z}_f , we have implemented (6) by choosing a different random split of the data set.

3.2. Meta-Testing

After meta-learning, during runtime, for a new frame f^{new} , we observe training channels $h_{l,f^{\text{new}}}$ with $l = 1, \dots, L^{\text{new}} + N + \delta - 1$. Accordingly, we have the training set $\mathcal{Z}_{f^{\text{new}}} = \{(h_{l,f^{\text{new}}}^N, h_{l+\delta,f^{\text{new}}})\}_{l=N}^{L^{\text{new}}+N-1}$. Based on these data and on the meta-trained hyperparameters $(\lambda^*, \bar{v}_{\lambda^*}^*)$, we train a channel predictor via ridge regression (1), obtaining

$$v_{f^{\text{new}}}^* = v_{\lambda^*}^*(\mathcal{Z}_{f^{\text{new}}} | \bar{v}_{\lambda^*}^*). \quad (7)$$

4. FAST TRAINING VIA ONLINE META-LEARNING

In this section, we briefly consider a scenario in which frames are transmitted in a streaming fashion so that, at any current frame f , we can use as meta-training data the channel data from past frames. Specifically, we assume that at frame f , we have access to data $\{\{h_{l,f'}\}_{l=1}^{L+N+\delta-1}\}_{f'=\min(f-M,1)}^{f-1}$ from the most recent M frames.

In order to facilitate an online update of the hyperparameter vector \bar{v} as more frames are collected, we develop an iterative gradient-based solution that follows EP [7].

To start, we define $L_{f'}(v) := \|X_{f'}^{\text{tr}} v - y_{f'}^{\text{tr}}\|^2$ for the training loss in frame f' ; $L_{f'}^{\text{inner}}(v | \bar{v}) := L_{f'}(v) + \lambda \|v - \bar{v}\|^2$ for the corresponding regularized loss; and $L_{f'}^{\text{outer}}(v) := \|X_{f'}^{\text{te}} v - y_{f'}^{\text{te}}\|^2$ for the corresponding test loss. With these definitions, the optimal hyperparameter vector (4) obtained at frame f based on the observation of data from frames $f' = \max(f-M, 1), \dots, f-1$ can be rewritten as

$$\bar{v}_{\lambda^*, M}^*(f) = \arg \min_{\bar{v}} \sum_{f'=\max(f-M, 1)}^{f-1} L_{f'}^{\text{outer}}(v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v})), \quad (8)$$

where $v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v}) = \arg \min L_{f'}^{\text{inner}}(v | \bar{v})$. EP tackles problem (8) via gradient descent [23] by introducing total loss function

$$L_{f'}^{\text{total}}(v | \bar{v}, \alpha) = L_{f'}^{\text{inner}}(v | \bar{v}) + \alpha L_{f'}^{\text{outer}}(v), \quad (9)$$

with an additional real-valued hyperparameter $\alpha \in \mathbb{R}$. Defining

$$\begin{aligned} v_{f'}^{\alpha, \bar{v}} &= \arg \min_v L_{f'}^{\text{total}}(v | \bar{v}, \alpha) \\ &= \arg \min_v \left\{ \|X_{f'}^{\text{tr}} v - y_{f'}^{\text{tr}}\|^2 + \alpha \|X_{f'}^{\text{te}} v - y_{f'}^{\text{te}}\|^2 + \lambda \|v - \bar{v}\|^2 \right\} \\ &= (A_{f'}^{\alpha})^{-1} \left((X_{f'}^{\alpha})^{\dagger} y_{f'}^{\alpha} + \lambda \bar{v} \right), \text{ with } A_{f'}^{\alpha} = ((X_{f'}^{\alpha})^{\dagger})^{\dagger} X_{f'}^{\alpha} + \lambda I, \end{aligned} \quad (10)$$

where the $L \times N$ matrix $X_{f'}^{\alpha} = [X_{f'}^{\text{tr}}, \sqrt{\alpha} X_{f'}^{\text{te}}]^{\top}$ stacks vertically the training input matrix $X_{f'}^{\text{tr}}$ and the weighted test input matrix $\sqrt{\alpha} X_{f'}^{\text{te}}$; and the $L \times 1$ vector $y_{f'}^{\alpha} = [y_{f'}^{\text{tr}}, \sqrt{\alpha} y_{f'}^{\text{te}}]^{\top}$ stacks vertically the training target vector $y_{f'}^{\text{tr}}$ and the weighted test target vector $\sqrt{\alpha} y_{f'}^{\text{te}}$. Based on definition (10), we have $v_{f'}^{\alpha=0, \bar{v}} = v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v})$, which is the solution of the inner problem in (8). Furthermore, it can be proved [7] that the gradient $\nabla_{\bar{v}} L_{f'}^{\text{outer}}(v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v}))$ of the objective of the meta-training problem (8) can be expressed as the limit

$$\begin{aligned} &\nabla_{\bar{v}} L_{f'}^{\text{outer}}(v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v})) \\ &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \left(\frac{\partial L_{f'}^{\text{total}}}{\partial \bar{v}}(v_{f'}^{\alpha, \bar{v}} | \bar{v}, \alpha) - \frac{\partial L_{f'}^{\text{total}}}{\partial \bar{v}}(v_{f'}^{0, \bar{v}} | \bar{v}, 0) \right) \\ &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \left(2\lambda (v_{f'}^{0, \bar{v}} - v_{f'}^{\alpha, \bar{v}}) \right). \end{aligned} \quad (11)$$

With a small enough scalar $\alpha \neq 0$, regardless of its sign, we can hence estimate the gradient $\nabla_{\bar{v}} L_{f'}^{\text{outer}}(v_{\lambda^*}^*(\mathcal{Z}_{f'}^{\text{tr}} | \bar{v}))$ as

$$\begin{aligned} \frac{1}{\alpha} (2\lambda (v_{f'}^{0, \bar{v}} - v_{f'}^{\alpha, \bar{v}})) &= \frac{1}{\alpha} \left(2\lambda \left((A_{f'}^{\alpha})^{-1} \left((X_{f'}^{\alpha})^{\dagger} y_{f'}^{\alpha} + \lambda \bar{v} \right) \right. \right. \\ &\quad \left. \left. - (A_{f'}^{\alpha})^{-1} \left((X_{f'}^{\alpha})^{\dagger} y_{f'}^{\alpha} + \lambda \bar{v} \right) \right) \right), \end{aligned} \quad (12)$$

where we have used (10). In order to obtain an online learning rule, we propose to apply gradient descent using gradient estimate (12) and adopting as initialization the final iterate produced at frame $f-1$.

5. EXPERIMENTS

In this section, we present a number of experiments for both offline and online meta-learning¹.

Random Doppler Frequency (Offline): Consider the Doppler spectra following Jakes' model

$$S_1(w) = \frac{1}{\frac{w_m}{2} \sqrt{1 - (w/w_m)^2}}, \quad |w| \leq w_m \quad (13)$$

and the rounded model [24]

$$\begin{aligned} S_2(w) &= \alpha \left[a_0 + a_2 (w/w_m)^2 + a_4 (w/w_m)^4 \right], \quad |w| \leq w_m, \\ \text{with } \alpha &= \frac{1}{\frac{w_m}{\pi} [a_0 + a_2/3 + a_4/5]}. \end{aligned} \quad (14)$$

We consider $a_0 = 1, a_2 = -1.72, a_4 = 0.785$ [24]. The Doppler spectrum $S_f(w)$ in each frame f is assumed to be generated by first picking at random, with the same probability, a Jakes' or rounded spectrum; and then selecting a

¹Code for regenerating the results of this paper can be found at <https://github.com/kclip/scalar-channel-meta-prediction>

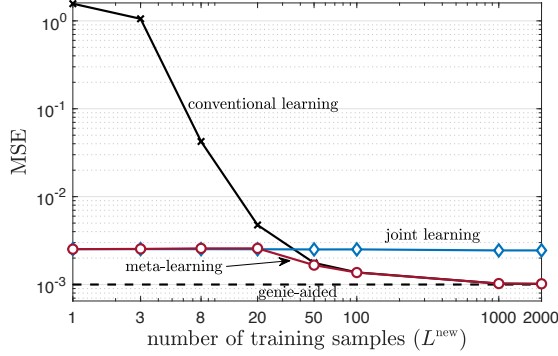


Fig. 2: Channel prediction performance as a function of the number of training samples L^{new} , with window size $N = 5$ and lag $\delta = 3$ for the setting with random maximum Doppler frequencies.

random Doppler frequency $w_m/2\pi \sim \mathcal{U}_{[0.05,0.1]}$. We fix $F = 80$ meta-training frames with number of slots per frame $L = L^{\text{tr}} + L^{\text{te}} = L^{\text{new}} + 100$.

We compare the proposed meta-learning scheme to (i) *conventional learning*, which directly applies the regressor (7) with $\lambda = 0$; (ii) *genie-aided prediction*, which implements LMMSE strategy based on knowledge of the Doppler spectrum; (iii) *joint learning*, which utilizes previous channel data at frames f jointly with the current channel at frame f^{new} to obtain the regressor $v_{f^{\text{new}}}^*$ by solving the problem

$$\min_v \left(\sum_f \sum_l \|v^\dagger h_{l,f}^N - h_{l+\delta,f}\|^2 + \sum_l \|v^\dagger h_{l,f^{\text{new}}}^N - h_{l+\delta,f^{\text{new}}}\|^2 \right). \quad (15)$$

Conventional learning and joint learning schemes correspond to LMMSE strategies based on covariance matrices estimated using pilots received, respectively, only in the new frame or in all the previous frames along with the new frame.

In Fig. 2, we compute the mean squared error (MSE) averaged over 10,000 samples. For prediction, we consider window size $N = 5$ with lag size $\delta = 3$. During meta-testing, we compute the channel prediction MSE with varying number of training samples L^{new} for the new frame f^{new} as shown in Fig. 2. All the considered schemes except for joint learning are seen to converge to the optimal performance as L^{new} grows, while meta-learning is observed to predict future channels even only with a single sample, i.e., with $L^{\text{new}} = 1$. Additionally, we also verified (not shown) that deep neural network-based predictors such as recurrent neural network or multilayer perceptron could not outperform the linear predictors considered in this work, as also reported in [5, 6].

Standard Channel Model (Offline): Here, we evaluate the proposed solution under the 5G standard spatical channel model (SCM) model [20]. To obtain frequency-flat fading, from the CDL-C channel model [20], we take only the first cluster into account to generate the channel model $h_{l,f}$. The Doppler frequency for each path is randomly scaled to have values between $w_{d,f}/2\pi \in [10, 100]$ by randomly changing the velocity of the user. Fig. 3 shows channel prediction performance




Fig. 3: Channel prediction performance as a function of the number of training samples L^{new} , with window size $N = 5$ and lag $\delta = 3$ under single clustered CDL-C model with random user velocity at each frame f [20].

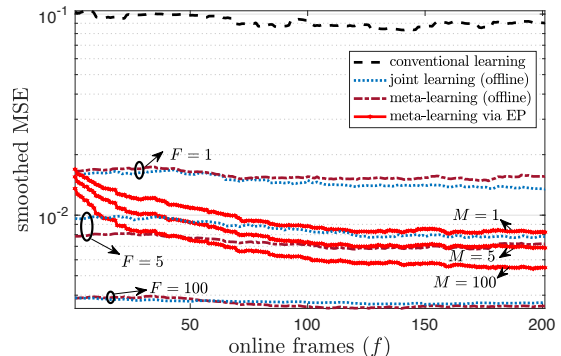


Fig. 4: Channel prediction performance as a function of current frame index f with window size $N = 5$ and lag $\delta = 3$ under single clustered CDL-C model with random user velocity for each frame f [20]. MSE is averaged over 10 independent experiments and smoothed over 100 frames.

with the same setting as Fig. 2, while considering 20 dB of channel estimation noise with $P = 100$ pilots. We observe that the proposed meta-learning approach works well even for channel models that do not strictly follow stationary Gaussian processes.

Gradient-Based Meta-Learning (Online): Lastly, we consider an online scenario in which frames are observed in a streaming fashion. We assume same standard channel model as in Sec. 5 during frames $f = 1, \dots, 300$ and compute the MSE at each frame f with fixed lambda $\lambda = 1$. We average over 10 independent experiments and plot the MSE smoothed with moving average with length 100. Fig. 4 shows that online meta-learning via EP outperforms conventional learning and offline joint learning even with $M = 1$, with the MSE decreasing as M is increased.

6. CONCLUSIONS

This paper considered linear channel prediction via meta-learning for both offline and online scenarios. The proposed meta-learned linear filters via closed-form expression and EP, were seen to effectively reduce the pilot overhead needed in each frame. Future work may consider prediction for frequency-selective MIMO fading channel.

7. REFERENCES

- [1] Mate Boban, Marco Giordani, and Michele Zorzi, "Predictive quality of service (PQoS): The next frontier for fully autonomous systems," *arXiv preprint arXiv:2109.09376*, 2021.
- [2] Peter Hoehner, Stefan Kaiser, and Patrick Robertson, "Two-dimensional pilot-symbol-aided channel estimation by Wiener filtering," in *Proc. ICASSP*, Munich, Germany, April 1997.
- [3] Alexandra Duel-Hallen, Shengquan Hu, and Hans Hallen, "Long-range prediction of fading signals," *IEEE Signal Processing Magazine*, vol. 17, no. 3, pp. 62–75, 2000.
- [4] Changkee Min, Namseok Chang, Jongsub Cha, and Joonhyuk Kang, "MIMO-OFDM downlink channel prediction for IEEE802.16e systems using Kalman filter," in *Proc. WCNC*. IEEE, 2007, pp. 942–946.
- [5] Hwanjin Kim *et al.*, "Massive MIMO channel prediction: Kalman filtering vs. machine learning," *IEEE Trans. Commun.*, vol. 69, no. 1, pp. 518–528, 2020.
- [6] Wei Jiang, Mathias Strufe, and Hans Dieter Schotten, "Long-range mimo channel prediction using recurrent neural networks," in *Proc. CCNC*, Las Vegas, USA, 2020.
- [7] Benjamin Scellier and Yoshua Bengio, "Equilibrium propagation: Bridging the gap between energy-based models and backpropagation," *Frontiers in Computational Neuroscience*, vol. 11, pp. 24, 2017.
- [8] Nicolas Zucchet, Simon Schug, Johannes von Oswald, Dominic Zhao, and Joao Sacramento, "A contrastive rule for meta-learning," *arXiv preprint arXiv:2104.01677*, 2021.
- [9] Sangwoo Park, Hyeryung Jang, Osvaldo Simeone, and Joonhyuk Kang, "Learning to demodulate from few pilots via offline and online meta-learning," *IEEE Trans. Signal Processing*, vol. 69, pp. 226–239, 2020.
- [10] Kfir M Cohen, Sangwoo Park, Osvaldo Simeone, and Shlomo Shamai, "Learning to learn to demodulate with uncertainty quantification via Bayesian meta-learning," *arXiv preprint arXiv:2108.00785*, 2021.
- [11] Tomer Raviv *et al.*, "Meta-viterbinet: Online meta-learned viterbi equalization for non-stationary channels," in *Proc. ICC workshops*, Montreal, Canada, June 2021.
- [12] Sangwoo Park, Osvaldo Simeone, and Joonhyuk Kang, "Meta-learning to communicate: Fast end-to-end training for fading channels," in *Proc. ICASSP*, Barcelona, Spain, May 2020.
- [13] Yihan Jiang, Hyeji Kim, Himanshu Asnani, and Sreeram Kannan, "Mind: Model independent neural decoder," in *Proc. SPAWC*, Cannes, France, July 2019.
- [14] Sangwoo Park, Osvaldo Simeone, and Joonhyuk Kang, "End-to-end fast training of communication links without a channel model via online meta-learning," in *Proc. SPAWC*, Atlanta, USA, May 2020.
- [15] Mathieu Goutay, Fayçal Ait Aoudia, and Jakob Hoydis, "Deep hypernetwork-based mimo detection," in *Proc. SPAWC*, Atlanta, USA, May 2020.
- [16] Yi Yuan, Gan Zheng, Kai-Kit Wong, Björn Ottersten, and Zhi-Quan Luo, "Transfer learning and meta learning-based fast downlink beamforming adaptation," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1742–1755, 2020.
- [17] Juping Zhang, Yi Yuan, Gan Zheng, Ioannis Krikidis, and Kai-Kit Wong, "Embedding model based fast meta learning for downlink beamforming adaptation," *IEEE Trans. Wireless Commun.*, 2021.
- [18] Ivana Nikoloska and Osvaldo Simeone, "Black-box and modular meta-learning for power control via random edge graph neural networks," *arXiv preprint arXiv:2108.13178*, 2021.
- [19] Robert G. Gallager, *Principles of digital communication*, Cambridge University Press, 2008.
- [20] 3GPP, "Study on channel model for frequencies from 0.5 to 100 ghz," *TR 38.901*, 2020.
- [21] Sharu Theresa Jose and Osvaldo Simeone, "Information-theoretic generalization bounds for meta-learning and applications," *Entropy*, vol. 23, no. 1, pp. 126, 2021.
- [22] Giulia Denevi, Carlo Ciliberto, Dimitris Stamos, and Massimiliano Pontil, "Learning to learn around a common mean," in *Proc. NIPS*, Montreal, Canada, Dec. 2018.
- [23] Osvaldo Simeone, "A brief introduction to machine learning for engineers," *Foundations and Trends in Signal Processing*, vol. 12, no. 3-4, pp. 200–431, 2018.
- [24] K. V. S. Hari *et al.*, "Channel models for fixed wireless applications," *IEEE 802.16 Broadband wireless access working group*, 2003.