

Unsupervised-Learning Power Control for Cell-Free Wireless Systems

Rasoul Nikbakht, Anders Jonsson, Angel Lozano

Universitat Pompeu Fabra (UPF)

08018 Barcelona, Spain.

Email: {rasoul.nikbakht, anders.jonsson, angel.lozano}@upf.edu

Abstract—This paper studies the viability of feedforward neural networks (NNs) for centralized power control in the uplink of cell-free wireless systems with matched-filter reception. The formulation relies only on large-scale channel behaviors as inputs, without the need for user location information, and on unsupervised learning, to avoid the onerous precomputation of training data that supervised learning would necessitate for every system or environment modification. Two different power control objectives are entertained, and for both of them the NN closely approximates the optimum solutions produced by convex solvers while vastly reducing the complexity, thereby opening the door to power control implementations for very large systems.

I. INTRODUCTION

Cell-free wireless systems can be interpreted as a deconstruction of cellular massive MIMO where, rather than concentrated at cell sites, the antennas are scattered, and the associations between users and cells are released. What results is a dense infrastructure of access points (APs), each featuring one or a few antennas, with every user potentially communicating with every AP [1]–[6]. This can also be seen as a form of network MIMO [7]–[9] or of centralized radio access [10], [11]. At the expense of extensive fronthaul, cell-free systems offer multiple advantages over their cellular counterparts, including large-scale diversity and reduced distances to the end users.

Focusing on the uplink, one of the challenges that arises is power control. Given the major differences in pathloss and shadowing among links, large-scale power control appears crucial to keep users near an AP from overwhelming other users' signals at that AP, and to avoid major performance disparities altogether.

The uplink power optimization in a cell-free system, and in fact in any multiple-access setting with linear reception, is often a nonconvex problem [12], [13]. Specifically, the maximization of the weighted sum of the user spectral efficiencies—a maximization that allows delimiting the spectral efficiency combinations that are fundamentally feasible—is in general nonconvex. This is because, while increasing a user's transmit power improves that user's reception, it also inflicts worse interference to the rest of users, and vice versa. Fortunately, less general yet highly relevant power control objectives, say the maximization of the minimum signal-to-interference-plus-noise ratio (SINR), can be cast in convex form [1, Sec. IV]. These objectives are hence particularly appealing, and in small systems a real-time convex solver could control the

transmit powers. However, this approach does not scale to hundreds or let alone thousands of users.

To scale up the cell-free power control process, we seek to exploit the ability of feedforward neural networks (NNs) to approximate arbitrary function mappings [14], an ability that is also being applied to tackle other problems in wireless communications [15]–[18]. Of particular relevance to this work is [18], which applies a NN with supervised learning to approximate the mapping between user locations and downlink transmit powers in a cellular massive-MIMO system. The features that we desire for our cell-free NN-based power control are somewhat different, precisely:

- Being based only on the channel gains between users and APs, with no reliance on the user positions. Such location information could be unreliable, or outright unavailable, especially in indoor settings.
- Tracking only the *large-scale* behavior of these channel gains, i.e., their local-averages. This is congruent with a centralized approach, as it averts having to recompute the transmit powers every few milliseconds and every few hundred kilohertz in reaction to small-scale fading fluctuations.
- Unsupervised learning. This avoids the precomputation of tailored training data for every conceivable environment and costly retraining every time there are changes in the system.

Given the above, the power control instrument of our choice is an unsupervised NN approximating the mapping between the system's large-scale channel gains and the uplink transmit powers with specific performance objectives.

II. SYSTEM AND CHANNEL MODELS

The systems under consideration feature N APs and K users per time-frequency resource unit, with N substantially larger than K so as to render matched-filter reception effective. APs and users are equipped with a single antenna, omnidirectional in azimuth.

Every user can potentially communicate with every AP on each time-frequency resource. A share of the resource units are reserved for pilot transmissions from the users, based on which the channels are estimated by the APs.

A. Large-scale Modeling

The APs and users are randomly placed, such that their locations conform to respective (mutually independent)

binomial point processes. As the system grows, these converge to Poisson point processes.

The operating bandwidth is in the microwave range. The signals are subject to pathloss with exponent η as well as log-normal shadowing, and the combination of both phenomena gives rise to a large-scale channel gain $G_{n,k}$ between the n th AP and the k th user.

Letting P denote a user's radiated power, measured at 1 m so that no scaling constants are needed, and with σ^2 the noise power, $\frac{P}{\sigma^2}$ is the SNR at 1 m. At the n th AP, the local-average SNR from user k is then $\text{SNR}_{n,k} = G_{n,k} \frac{P}{\sigma^2}$. The local-average SNRs, and the large-scale parameters in general, are stable and known.

B. Small-scale Modeling

Besides $G_{n,k}$, the channel that connects the k th user with the n th AP features a small-scale fading coefficient $h_{n,k} \sim \mathcal{N}_{\mathbb{C}}(0, 1)$, independent across users and APs.

C. Simulation Environment

To generate performance distributions over many system snapshots, we resort to a wrapped-around (i.e., without boundaries) universe. Under the assumption that the AP positions are agnostic to the radio propagation, shadow fading has been shown to render networks approximately Poisson-like from the perspective of any user [19]. This approximation sharpens as the shadowing variance grows, being highly precise for values of interest [19]–[21]. Capitalizing on this result, in our simulator the AP positions are drawn uniformly, avoiding the need for explicit modeling of the shadowing as it is then already implicitly captured by the geometry. Likewise, the user positions are drawn uniformly.

Results are generated from 10000 system snapshots with $N = 100$ APs each, ensuring a 95% confidence interval below 0.03 dB.

III. CELL-FREE UPLINK FORMULATION

A. Channel Estimation

Disregarding pilot contamination, which can be kept at bay through procedures such as the ones described in [1, Sec. IV] or in [22], [23], the linear MMSE estimate $\hat{h}_{n,k}$ gathered by the system upon observation of a pilot transmission per user and coherence block satisfies $h_{n,k} = \hat{h}_{n,k} + \tilde{h}_{n,k}$ where [13, Sec. 3.7]

$$\tilde{h}_{n,k} \sim \mathcal{N}_{\mathbb{C}}\left(0, \frac{1}{1 + \text{SNR}_{n,k}}\right) \quad (1)$$

is uncorrelated error and the channel estimate's power equals

$$\mathbb{E}[|\hat{h}_{n,k}|^2] = \frac{\text{SNR}_{n,k}}{1 + \text{SNR}_{n,k}}. \quad (2)$$

B. Data Transmission

Upon payload data transmission, the n th AP observes

$$y_n = \sum_{k=0}^{K-1} \sqrt{G_{n,k}} h_{n,k} \sqrt{p_k} s_k + v_n, \quad (3)$$

where s_k is the unit-variance symbol transmitted by user k while $p_k \in [0, 1]$ is its power control coefficient. For its part, $v_n \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2)$ is the noise.

With matched filtering, the signal of user k is recovered as $\sum_{n=0}^{N-1} w_{n,k}^* y_n$ where $w_{n,k} = \sqrt{G_{n,k}} \hat{h}_{n,k}$. This gives

$$\begin{aligned} \sum_{n=0}^{N-1} w_{n,k}^* y_n &= \underbrace{\sum_{n=0}^{N-1} G_{n,k} \sqrt{p_k} \hat{h}_{n,k}^* \hat{h}_{n,k} s_k}_{\text{Desired Signal: } S_k} \\ &+ \underbrace{\sum_{n=0}^{N-1} \sqrt{G_{n,k}} \hat{h}_{n,k}^* \sum_{\ell \neq k} \sqrt{G_{n,\ell}} \sqrt{p_\ell} \hat{h}_{n,\ell} s_\ell}_{\text{Multiuser Interference: } I_k} \\ &+ \underbrace{\sum_{n=0}^{N-1} \sqrt{G_{n,k}} \hat{h}_{n,k}^* \sum_{\ell=0}^{K-1} \sqrt{G_{n,\ell}} \sqrt{p_\ell} \tilde{h}_{n,\ell} s_\ell}_{\text{Channel Estimation Errors: } E_k} \\ &+ \underbrace{\sum_{n=0}^{N-1} \sqrt{G_{n,k}} \hat{h}_{n,k}^* v_n}_{\text{Filtered Noise: } V_k}, \end{aligned} \quad (4)$$

from which the SINR of user k , conditioned on the known $\{\hat{h}_{n,k}\}_{n=0}^{N-1}$, equals

$$\text{sinr}_k = \frac{\mathbb{E}[|S_k|^2 | \{\hat{h}_{n,k}\}_{n=0}^{N-1}]}{\mathbb{E}[|I_k|^2 + |E_k|^2 + |V_k|^2 | \{\hat{h}_{n,k}\}_{n=0}^{N-1}]} \quad (5)$$

$$= \frac{P \left(\sum_{n=0}^{N-1} G_{n,k} \sqrt{p_k} |\hat{h}_{n,k}|^2 \right)^2}{\text{den}_k} \quad (6)$$

with

$$\begin{aligned} \text{den}_k &= P \sum_{\ell \neq k} p_\ell \left| \sum_{n=0}^{N-1} \sqrt{G_{n,k}} \sqrt{G_{n,\ell}} \hat{h}_{n,k}^* \hat{h}_{n,\ell} \right|^2 \\ &+ P \sum_{\ell=0}^{K-1} p_\ell \sum_{n=0}^{N-1} G_{n,k} G_{n,\ell} |\hat{h}_{n,k}|^2 \text{MMSE}_{n,\ell} \\ &+ \sum_{n=0}^{N-1} G_{n,k} |\hat{h}_{n,k}|^2 \sigma^2. \end{aligned} \quad (7)$$

After some simplifications, and substituting for $\text{MMSE}_{n,k}$,

$$\text{sinr}_k = \frac{p_k \left(\sum_{n=0}^{N-1} \text{SNR}_{n,k} |\hat{h}_{n,k}|^2 \right)^2}{\text{den}_k} \quad (8)$$

with

$$\begin{aligned} \text{den}_k &= \sum_{\ell \neq k} p_\ell \left| \sum_{n=0}^{N-1} \sqrt{\text{SNR}_{n,k} \text{SNR}_{n,\ell}} \hat{h}_{n,k}^* \hat{h}_{n,\ell} \right|^2 \\ &+ \sum_{\ell=0}^{K-1} p_\ell \sum_{n=0}^{N-1} \frac{\text{SNR}_{n,k} \text{SNR}_{n,\ell}}{1 + \text{SNR}_{n,\ell}} |\hat{h}_{n,k}|^2 \\ &+ \sum_{n=0}^{N-1} \text{SNR}_{n,k} |\hat{h}_{n,k}|^2. \end{aligned} \quad (9)$$

In interference-limited conditions, (8)–(13) simplify into

$$\text{sir}_k = \frac{p_k \left(\sum_{n=0}^{N-1} G_{n,k} |h_{n,k}|^2 \right)^2}{\sum_{\ell \neq k} p_\ell \left| \sum_{n=0}^{N-1} \sqrt{G_{n,k} G_{n,\ell}} h_{n,k}^* h_{n,\ell} \right|^2}. \quad (10)$$

IV. SIR-BASED POWER CONTROL OBJECTIVES

If, in (10), we replace $|h_{n,k}|^2$ and $|h_{n,\ell}|^2$ by their expected value (unity), and we replace the cross-terms containing $h_{n,k}^* h_{n,\ell}$ with $k \neq \ell$ also by their expected value (zero), we obtain

$$\text{SIR}_k = \frac{p_k \left(\sum_{n=0}^{N-1} G_{n,k} \right)^2}{\sum_{\ell \neq k} p_\ell \sum_{n=0}^{N-1} G_{n,k} G_{n,\ell}}, \quad (11)$$

which, except for the absence of a self-interference term corresponding to $\ell = k$, coincides with the SIR that would be achieved by a receiver that, rather than decode the signal of user k based on $\{h_{n,k}\}_{n=0}^{N-1}$, exploited channel hardening [1]. From $\{\text{SIR}_k\}_{k=0}^{K-1}$, certain power control objectives can be formulated as geometric programming problems and tackled with off-the-shelf convex optimization toolboxes. The power control coefficients thereby obtained can then be plugged into (10) to assess the $\{\text{sir}_k\}_{k=0}^{K-1}$ that receivers reliant on channel estimates would attain.

To date, cell-free power control has almost exclusively pursued a max-min objective, hence such is our starting point as well, but our scope expands to include a second and highly relevant objective, namely max-product.

A. (Soft) Max-Min

This pursuit of this objective, or more precisely a soft version thereof, can be formulated as the minimization over $\{p_k\}_{k=0}^{K-1}$ of the cost

$$L_{\text{MM}} = \sum_{k=0}^{K-1} e^{-\alpha_k \text{SIR}_k} \quad (12)$$

where, recall, $\{\text{SIR}_k\}_{k=0}^{K-1}$ are functions of $\{p_k\}_{k=0}^{K-1}$ as per (11), while $\{\alpha_k\}_{k=0}^{K-1}$ are parameters that determine the softness of the max-min operation. As these parameters grow large, the cost becomes dominated by the smallest SIR, and the minimization of L_{MM} converges to a hard max-min formulation. Conversely, if the parameters decrease, the behavior softens as SIRs other than the smallest begin playing a role; this prevents users in highly adverse situations from dragging down the entire system.

B. Max-Product

The maximization of $\prod_{k=0}^{K-1} \text{SIR}_k$, or more precisely its logarithm, can be posed as the minimization over $\{p_k\}_{k=0}^{K-1}$ of the cost

$$L_{\text{MP}} = -\frac{1}{K} \sum_{k=0}^{K-1} \beta_k \log \text{SIR}_k \quad (13)$$

where $\{\beta_k\}_{k=0}^{K-1}$ are again parameters. When these parameters are equal, the combination $\{\text{SIR}_k\}_{k=0}^{K-1}$ minimizing

L_{MP} is termed the *proportional fair* operating point, which exhibits very satisfying properties in terms of the tradeoff between aggregate performance and fairness [24]. The parameters provide further freedom to prioritize users.

V. UNSUPERVISED LEARNING PROCEDURE

As mentioned, our interest is in an unsupervised feed-forward NN that accepts as inputs the large-scale channel gains and outputs the power control coefficients $\{p_k\}_{k=0}^{K-1}$. While, in supervised learning, the cost would be a function of the difference between the predicted and the correct output, in our unsupervised approach we never explicitly represent the correct output for each input. Rather, we define the cost as L_{MM} or L_{MP} and compute the gradient directly on it. We do not know the correct output for a given input, but we can compute the cost in L_{MM} or L_{MP} given the current prediction of $\{p_k\}_{k=0}^{K-1}$ and attempt to update such prediction in order to minimize this cost.

Recognizing that, with matched-filter reception, the effective gain between user k and the system is

$$G_k = \sum_{n=0}^{N-1} G_{n,k}, \quad (14)$$

rather than feed the NN with the NK large-scale gains, we reduce the input dimensionality and feed only the effective gains $\{G_k\}_{k=0}^{K-1}$. This compresses the learning process rather drastically while affording excellent performance, as exemplified in Section VI.

The effective gains $\{G_k\}_{k=0}^{K-1}$ are further preprocessed by converting them to log-scale, subtracting their mean, and dividing by their standard deviation, following which they are fed to an input layer equipped with rectified linear unit (RLU) activation functions. After feature extraction on the part of this input layer, a hidden layer processes the data also via RLU activation functions, and an output layer with linear activation functions spits out power control coefficients in log-scale; this guarantees positive and nonzero outputs and prevents numerical problems.

From the NN outputs and the corresponding large-scale gains, the cost of choice (soft max-min or max-product) is quantified and an Adam optimizer—a typical algorithm to update NN weights iteratively [25]—of learning rate 0.001 is applied to minimize such cost. To avoid overfitting, L2-norm regularization is used in conjunction with the Adam optimizer. Specifically, a portion $\lambda = 0.001$ of the L2 norm of the weights is added to the cost.

The complete scheme is illustrated in Fig. 1, while the parameters of the NN are summarized in Table I.

To simplify the learning, rather than a single large database we generate multiple small databases. Precisely, M databases of 6400 system snapshots are produced and, over each one, L updates of the NN weights take place; each such update uses a randomly selected batch of $B = 64$ snapshots. In total, $6400M$ system snapshots are produced for learning purposes, and the NN weights are updated LM times.

The learning process is a nonconvex problem, and the weights of the NN are initialized randomly, hence a single

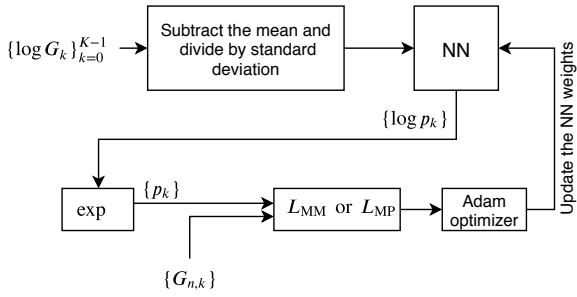


Fig. 1: NN learning procedure.

TABLE I: NN parameters.

	Input layer	Hidden layer	Output layer
Number of neurons	2000	200	40
Activation function	RLU	RLU	Linear
Regularization	L2 norm $\lambda = .001$	L2 norm $\lambda = .001$	L2 norm $\lambda = .001$

NN does not provide sufficient effectiveness guarantees. To address this issue, we train 100 distinct NNs and the performance is averaged over them.

VI. PERFORMANCE EVALUATION

We consider a system with $N = 100$ and $K = 40$, and with a pathloss exponent $\eta = 3.8$. (Much larger systems can be handled by the NN, the bottleneck are the baseline convex optimizations.) Both the NN and the convex optimization are driven by the SIRs in (11) while the performance evaluations are conducted with the SIRs in (10), averaged over the small-scale fading.

A. (Soft) Max-Min

Shown in Fig. 2 is the learning curve of the NN with this objective and $\alpha_k = 1$ for $k = 0, \dots, K-1$. Since the learning occurs over small databases, there is a certain degree of overfitting. A spike ensues whenever a new database is fed in, confirming that the NN was partially overfitted to the previous database. However, the amplitude of these spikes decreases gradually as more data is fed in; the overfitting vanishes and the NN can then generalize well to unseen data. We note that the cost need not vanish, as the optimum power allocation in general does not map to a zero cost.

Fig. 3 portrays the CDF of sir_k achieved by the NN-based power control under a soft max-min objective with $\alpha_k = 1$ for $k = 0, \dots, K-1$, as well as its convex-optimization counterpart and the hard max-min solution obtained also via convex optimization. This figure prompts two observations:

- The NN matches very closely its convex optimization counterpart (with orders-of-magnitude less complexity, as established later in the section).
- A hard max-min approach is unwise, as its solution is dragged down by worst-case users. In contrast, a soft max-min allows those users to underperform while the rest achieve substantially higher SIRs.

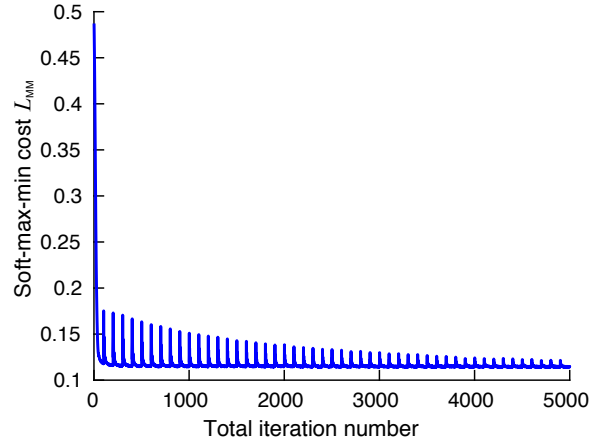


Fig. 2: Learning curve (averaged over 100 NNs) for soft max-min with unit parameters.

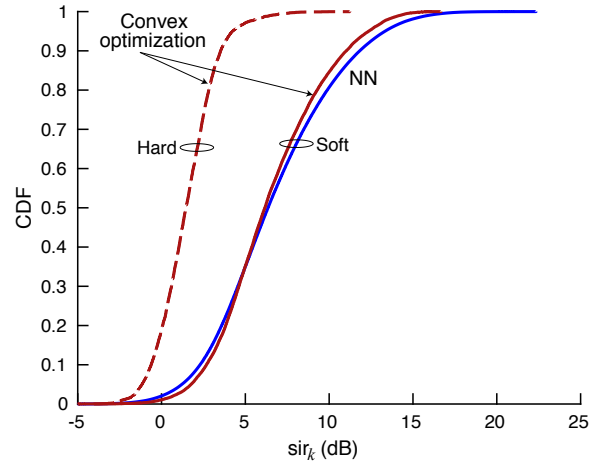


Fig. 3: CDF of sir_k : NN and convex optimization for soft max-min with unit parameters, and convex optimization for hard max-min.

The superior performance of softer max-min versions, with $\alpha_k \approx 1$, is a welcome development for another reason, namely that the NN runs into numerical problems and becomes unstable as $\{\alpha_k\}_{k=0}^{K-1}$ grow large.

B. Max-Product

The learning curve for this objective is similar to its max-min brethren, hence for the sake of compactness it is not shown. The CDF of sir_k for max-product power control with the parameters β_k equal for all users is presented in the main plot of Fig. 4. The match between the NN and its convex-optimization counterpart is slightly less precise than for max-min, but still satisfactory. The figure's inset depicts the max-product itself, evidencing that the small discrepancy between the NN and the convex-optimization solutions is the manifestation of an equally small shortfall (about 0.5 dB) in the achieved objective.

C. Complexity

As a proxy for complexity, we invoke the running time on a common computational platform. Table II summarizes the average running time for both the NN and the CVX convex solver [26]. Once the NN has undergone

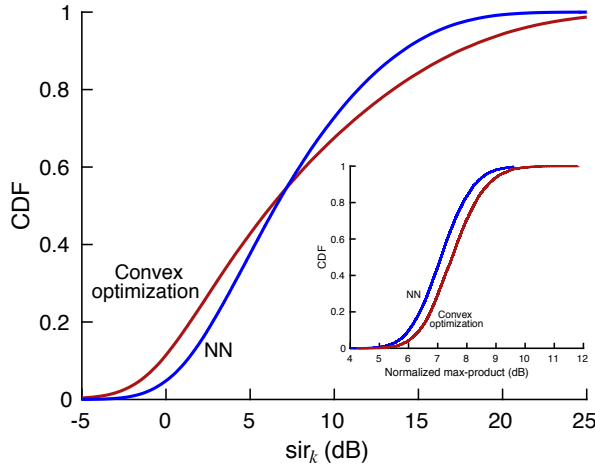


Fig. 4: Main plot: CDF of sir_k for the NN and the convex optimization power controls with max-product and equal parameters. Inset: corresponding CDFs of $\frac{1}{K} \sum_{k=0}^{K-1} 10 \log_{10} \text{sir}_k$.

TABLE II: Average running time (s) per system snapshot.

	Learning	Test time
NN	1230	< 0.01
Max-min CVX	-	2.4
Max-product CVX	-	2.7

the learning process, its evaluation for a given system snapshot is as fast as a matrix multiplication while the convex optimization process requires well over two orders of magnitude more time.

The training time is not a primary concern because retraining needs to take place only sporadically, upon substantial changes in the system or the environment.

VII. SUMMARY

This paper has shown that a feedforward NN with unsupervised learning can closely approximate the performance of vastly more computationally demanding convex solvers, opening the door to centralized power control for very large cell-free systems. Such centralized control can ensure superior performance while enforcing specific SIR objectives such as the max-min and max-product considered herein.

In our implementation, the NN is only fed the K effective channel gains $\{G_k\}$. We have verified that, if the NK large-scale channel gains $\{G_{n,k}\}$ themselves are used as inputs, the approximation tightens even further, at the expense of a longer learning stage.

ACKNOWLEDGMENT

This work was supported by the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502) and by the European Research Council under the H2020 Framework Programme/ERC grant agreement 694974.

REFERENCES

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

[2] G. Interdonato, H. Q. Ngo, E. G. Larsson, and P. Frenger, "How much do downlink pilots improve cell-free massive MIMO?" *IEEE Global Commun. Conf. (GLOBECOM'16)*, pp. 1–7, 2016.

[3] E. Nayebi, A. Ashikhmin, T. L. Marzetta, H. Yang, and B. D. Rao, "Precoding and power optimization in cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 16, pp. 4445–4459, Jul. 2017.

[4] H. Q. Ngo, L.-N. Tran, T. Q. Duong, M. Matthaiou, and E. G. Larsson, "On the total energy efficiency of cell-free massive MIMO," *IEEE Trans. Green Commun. Networking*, vol. 2, no. 1, pp. 25–39, Mar. 2018.

[5] M. Attarifar, A. Abbasfar, and A. Lozano, "Modified conjugate beamforming for cell-free massive MIMO," *IEEE Wireless Commun. Letters*, vol. 9, 2019.

[6] A. Lozano and R. Nikbakht, "Uplink fractional power control for cell-free wireless networks," in *IEEE Int'l Conf. on Communications (ICC'19)*, May 2019.

[7] M. K. Karakayali, G. J. Foschini, and R. A. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Commun. Mag.*, vol. 3, no. 14, pp. 56–61, Aug. 2006.

[8] S. Venkatesan, A. Lozano, and R. Valenzuela, "Network MIMO: Overcoming intercell interference in indoor wireless systems," *Asilomar Conf. Signals, Systems and Computers*, pp. 83–87, 2007.

[9] S. Venkatesan, H. Huang, A. Lozano, and R. Valenzuela, "A WiMAX-based implementation of network MIMO for indoor wireless systems," *EURASIP J. Adv. Signal Processing*, vol. 2009, Sep. 2009.

[10] A. Checko *et al.*, "Cloud RAN for mobile networks – a technology overview," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 405–426, 1st Quart. 2015.

[11] S. Perlman and A. Forenza, "An introduction to pCell," Artemis Networks LLC, White paper, Tech. Rep., Feb. 2015. [Online]. Available: <http://www.rearden.com/artemis/An-Introduction-to-pCell-White-Paper-150224.pdf>

[12] M. Chiang, C. W. Tan, D. P. Palomar, D. O'Neill, and D. Julian, "Power control by geometric programming," *IEEE Trans. on Wireless Communications*, vol. 6, no. 7, pp. 2640–2651, 2007.

[13] R. W. Heath Jr. and A. Lozano, *Foundations of MIMO Communication*. Cambridge University Press, 2018.

[14] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.

[15] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Trans. Cognitive Commun. and Networking*, vol. 3, no. 4, pp. 563–575, 2017.

[16] H. Sun, X. Chen, Q. Shi, M. Hong, X. Fu, and N. D. Sidiropoulos, "Learning to optimize: Training deep neural networks for interference management," *IEEE Trans. Signal Processing*, vol. 66, no. 20, pp. 5438–5453, 2018.

[17] W. Cui, K. Shen, and W. Yu, "Spatial deep learning for wireless scheduling," *arXiv preprint arXiv:1808.01486*, 2018.

[18] L. Sanguinetti, A. Zappone, and M. Debbah, "Deep learning power allocation in massive MIMO," *arXiv preprint arXiv:1812.03640*, 2018.

[19] B. Błaszczyszyn, M. K. Karray, and H. P. Keeler, "Wireless networks appear Poissonian due to strong shadowing," *IEEE Trans. Wireless Communications*, vol. 14, pp. 4379–4390, Aug. 2015.

[20] G. George, R. K. Mungara, A. Lozano, and M. Haenggi, "Ergodic spectral efficiency in MIMO cellular networks," *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 2835–2849, May 2017.

[21] G. George, A. Lozano, and M. Haenggi, "Massive MIMO forward link analysis for cellular networks," *IEEE Trans. Wireless Commun.*, vol. 18, pp. 2964–2976, June 2019.

[22] O. Y. Bursalioglu, C. Wang, H. Papadopoulos, and G. Caire, "RRH based massive MIMO with 'on the fly' pilot contamination control," in *IEEE Int'l Conf. Commun. (ICC'16)*, 2016, pp. 1–7.

[23] M. Attarifar, A. Abbasfar, and A. Lozano, "Random vs structured pilot assignment in cell-free massive MIMO wireless networks," *IEEE Int'l Conf. Commun. Workshops (ICCW'18)*, May 2018.

[24] F. Kelly, "Charging and rate control for elastic traffic," *European Trans. on Telecommunications*, vol. 8, no. 1, pp. 33–37, 1997.

[25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[26] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," <http://cvxr.com/cvx>, Mar. 2014.