

SYMMeTRy: Exploiting MIMO Self-Similarity for Under-Determined Modulation Recognition

Wei Xiong, Lin Zhang, Maxwell McNeil, Petko Bogdanov, Mariya Zheleva
 Department of Computer Science, University at Albany – SUNY
 {wxiong, lzhang22, mmcneil2, pbogdanov, mzheleva}@albany.edu

Abstract—Modulation recognition (modrec) seeks to identify the modulation of a transmitter from corresponding spectrum scans. It is an essential functional component of future spectrum sensing with critical applications in dynamic spectrum access and spectrum enforcement. While predominantly studied in single-input single-output (SISO) systems, practical modrec for multiple-input multiple-output (MIMO) communications requires more research attention. Existing MIMO modrec impose stringent requirements of fully- or over-determined sensing front-end, i.e. the number of sensor antennas should exceed that at the transmitter. This poses a prohibitive sensor cost even for simple 2x2 MIMO systems and will severely hamper progress in flexible spectrum access. We design a MIMO modrec framework that enables efficient and cost-effective modulation classification for under-determined settings involving fewer sensor antennas than those used for transmission. Our key idea is to exploit the inherent multi-scale self-similarity of MIMO modulation IQ constellations, which persists in under-determined settings. Our framework, called **SYMMeTRy** (Self-similarity for MIMO Modulation Recognition), designs domain-aware classification features with high discriminative potential by summarizing regularities of symbol co-location in the MIMO constellation. To this end, we summarize the fractal geometry of observed samples to extract discriminative features for supervised MIMO modrec. We evaluate **SYMMeTRy** in a realistic simulation and in a small-scale MIMO testbed. We demonstrate that it maintains high and consistent performance across various noise regimes, channel fading conditions and with increasing MIMO transmitter complexity. Our efforts highlight **SYMMeTRy**'s high potential to enable efficient and practical MIMO modrec in spectrum sensing infrastructures with mixed-complexity sensors.

Index Terms—MIMO modulation recognition, under-determined spectrum sensing, machine learning, fractal dimension features, SVM



1 INTRODUCTION

Dynamic Spectrum Access (DSA) is projected as a key capability in 5G mobile networks, seeking to address the shortage of radio resources by opportunistic frequency reuse. While DSA promises improved network performance, it hinges on robust and affordable spectrum measurement in support of technology, policy and enforcement. For example, future spectrum enforcement [3] will need to automatically scan and track rogue transmitters without prior knowledge of their technology. Furthermore, emerging spectrum coexistence scenarios, such as in the CBRS and 6GHz bands [38], will require proactive and predictive transmitter sensing and profiling [39] to inform resource allocation [40]. An essential component of a transmitter's fingerprint is its modulation. This brings automatic modulation recognition (modrec) in the research spotlight [7], [33]. Traditionally, modrec was designed for cooperative *signal decoding*, and assumed prior knowledge of transmitter characteristics (e.g. signal shaping, bandwidth, central frequency and antenna complexity). However, emerging spectrum sensing systems [1], [24], [37], [44] that perform blind, sweep-based sensing, and the corresponding applications [3], [38], [40] that cannot hold prior assumptions of transmitter technology, inherently cannot cater to such stringent sensing requirements. Thus, in order to enable *emerging spectrum applications* we have to design novel modrec approaches for spectrum analytics without prior transmitter information.

While traditionally modrec has been tackled in the SISO context [2], [8], [12], [13], [22], [29], [41], [48], the ubiquity

of MIMO technology requires the design of robust and cost-efficient MIMO modrec with practical applicability to emerging spectrum sensing platforms [1], [24], [37]. The problem of MIMO modrec is more challenging than the SISO case since (i) the number of symbols in the IQ constellation grows exponentially with the number of transmit antennas and (ii) the channel state parameters grow quadratically with the number of antennas. To overcome these challenges, MIMO modrec requires the sensing infrastructure to be either fully- or over-determined, that is the number of antennas on the sensor should be equal or double that of the target transmitter [5], [18], [27], [51]. These requirements pose prohibitive cost for an individual sensor and present a major road-block to affordable and ubiquitous spectrum sensing. In addition, fully- and over-determined sensing generates large volumes of spectrum data posing high bandwidth, storage and computation requirements and further hampering spectrum analytics at scale.

To put this into perspective, Tbl. 1 presents a breakdown of the cost and data footprint of spectrum sensing when using a USRP B210 sensor at \$1,216 per board, collecting a short 10-second scan at 1MSps. A USRP B210 can support up to 2x2 MIMO operation with its embedded capabilities. We assume the use of an Octoclock-G CDA-2990 at \$1,927 for higher order MIMO setups. It is evident that both the monetary cost and the data footprint scale super-linearly with the MIMO complexity. For the worst case of over-determined sensing of a 4x4 MIMO setup the cost for a single sensor is \$11,655 generating 80GB of data for 10 sec-

	Cost, (\$)			Storage, (GB)		
	U	F	O	U	F	O
SISO	1,216	1,216	2,432	10	10	20
MIMO 2x2	1,216	2,432	6,791	10	20	40
MIMO 4x4	1,216	6,791	11,655	10	40	80

TABLE 1: Monetary cost and data footprint of a 10-second MIMO trace with under- (U), fully- (F) and over-determined (O) sensing using a USRP B210 at a sampling rate of 1MSps and an Octoclock-G.




Fig. 1: Illustration of MIMO self-similarity for a 2x2 QPSK MIMO signal.

onds. These costs are just for the radio hardware and would further be amplified if the host computer is accounted for.

In order to address this prohibitive cost we propose to utilize *under-determined spectrum sensing*, whereby a multi-antenna transmitter is scanned by a sensor with fewer than the transmitter’s antennas. While this has the potential to keep the monetary and data cost of MIMO sensing in check, it raises fundamental challenges for characterization. Existing MIMO modrec [5], [18], [27], [51] requires fully- or over-determined sensing as it typically relies on invariant statistical properties of the sample constellations. Thus, standard MIMO modrec cannot be readily-applied in the under-determined scenario. To address this, we propose a MIMO modrec framework called **SYMMeTRY** (Self-similarity for MIMO Modulation Recognition) that enables robust recognition for under-determined sensing. **SYMMeTRY** explores the self-similar geometric patterns in the MIMO constellation to extract domain-informed classification features. The key insight behind **SYMMeTRY** is that MIMO constellations exhibits self-similarity at different scales, which is well-preserved in under-determined settings and can be employed as a modrec feature. Fig. 1 illustrates the concept of self-similarity via an example 2x2 MIMO QPSK modulated signal.¹ A SISO QPSK constellation (top left corner of Fig. 1) contains a group of four clusters, one for each of the QPSK symbols. In the 2x2 MIMO QPSK case (bottom left), four new QPSK symbol groups “hatch” around each of the existing SISO positions. Therefore, for 2x2 QPSK MIMO there are at most 4^2 constellation symbols. Such templated replication behavior leads to a fractal-like self-similar organization of the MIMO constellation. The 2x2 QPSK MIMO constellation can, thus, be thought of as a 2-tier hierarchical representation (Fig. 1(right)), whereby at the first level, we observe the actual constellation composed of four SISO QPSK groups, while at the second level, we observe a meta-constellation that is determined by the centroids of each group at layer 1 and resembles a scaled version of a SISO QPSK constellation. This self-similar multi-scale organization can be observed for higher order modulations, whereby, the number of meta-constellation levels increases

1. For clarity, most of our examples focus on simple settings, such as 2x1 MIMO sensing of QPSK modulation. We note that our evaluation explores higher order modulations and more complex MIMO systems.




Fig. 2: System model for under-determined sensing. The right pane illustrates a 2x2 QPSK MIMO system. Existing MIMO modulation recognition requires fully- or over-determined sensing, meaning that the sensor antennas have to be equal or double these of the transmitter. We conceptualize under-determined sensing, as illustrated in the left pane of the figure, whereby a MIMO link is sensed with fewer antennas than these at the transmitter.

to $\log(M) - 1$ for increasing modulation order up to M . Furthermore, this self-similarity persists even when a scan is under-determined and is at the center of our design of modulation classification features. To capture the relative co-location of symbols within the constellation we extend tools from fractal geometry [11], [21]. We combine features based on the above patterns with higher order cumulants [14], [41] in our overall feature design.

Our paper makes the following key contributions:

- We conceptualize under-determined MIMO spectrum sensing and modulation classification.
- We are the first to formalize and utilize the self-similarity of MIMO constellations into an adaptive framework called **SYMMeTRY**, for robust feature design in support of under-determined MIMO modrec.
- We investigate the effects of the MIMO channel on the constellation self-similarity and recognition accuracy and show that **SYMMeTRY** is robust to changing channel conditions.
- Using realistic simulation and a testbed, we show that **SYMMeTRY** maintains high and consistent performance across various channel conditions, training scenarios and with increasing MIMO complexity.
- We evaluate **SYMMeTRY**’s runtime performance, demonstrating that the combination of feature extraction and classifier training takes in the order of several seconds (depending on instance size). Thus, **SYMMeTRY** has potential to support near-real-time modrec applications.

2 BACKGROUND

We introduce the general MIMO signal model and under-determined MIMO sensing. We detail the geometric self-similarity of MIMO constellations to provide intuition behind our methodology. We also analyze the relationship between MIMO channel properties such as (i) signal to noise ratio (SNR), (ii) coherence and (iii) pairwise antenna gain, and the geometric properties of the observed constellation.

2.1 MIMO signal and system models

The MIMO system is composed of N_t transmitting antennas and N_r receiving antennas [6], [26]. Fig. 2 (right) illustrates a 2x2 MIMO link (i.e. $N_r = 2$ and $N_t = 2$) using QPSK modulation. At any given instant, the signals $r(n) \in \mathbb{C}^{N_r}$ on the receiving side can be represented as a linear combination of the transmitted baseband signals $s(n) \in \mathbb{C}^{N_t}$ and instantaneous additive Gaussian noise $w(n) \in \mathbb{C}^{N_r}$:

$$r(n) = H s(n) + w(n), \quad (1)$$

where $H = [h_{ij}] \in \mathbb{C}^{N_r \times N_t}$ is the channel matrix specifying the pairwise channel response h_{ij} for each transmit-receive antenna pair (i, j) . The channel noise is modeled as sampled from a 0-mean normal distribution $w(n) \in \mathcal{N}(0, \sigma^2 I_{N_r})$ with 0 co-variance terms, i.e., $E[w(n)w(n)^H] = \sigma^2 I_{N_r}$, where I_{N_r} is the identity matrix of size N_r and $w(n)^H$ denotes the conjugate transpose of the row vector $w(n)$.

Depending on how application data is divided over the individual transmitter streams N_t and how each stream is modulated, there are several different ways to realize a MIMO transmitter [34]. In terms of application data assignment we have (i) MIMO with *spatial diversity* and (ii) MIMO with *spatial multiplexing*. The former passes redundant application data on each transmitter antenna allowing the receiver to decode the minimal-error stream. This redundancy leads to lower throughput in comparison with spatial multiplexing, but enables high probability of successful data decoding at the receiver. In *spatial multiplexing* MIMO each transmitter stream handles a unique portion of the application data resulting in increased throughput at the cost of higher sensitivity to poor channel conditions.

We also differentiate between *direct-mapped* and *precoded* MIMO, based on the power and modulation assignment across transmit streams. The former allocates the same power and modulation to all streams without considering the channel conditions, while the latter performs channel estimation and adaptively assigns power and modulation to each stream to maximize throughput. Thus, the choice between direct-mapping and precoding offers a trade-off between throughput and implementation complexity.

In this paper, we explore MIMO modulation recognition in all of the above MIMO realizations.

2.2 Limitations of existing MIMO modrec approaches

Modulation recognition can be viewed as a classification task: given a set of N IQ samples (an “instance”) the goal is to determine the modulation from which the observations are sampled. Classifiers used in such tasks are trained in a supervised manner, i.e. they require annotated instances in order to learn to recognize modulations (called classes). Most existing MIMO modrec algorithms require prior knowledge of the channel conditions [5], [16], [18], [27], [30]. In the context of the MIMO signal model (Eq. (1)), this means that existing approaches require prior knowledge of the channel response matrix H . Once this information is available, the paths are considered individually and common features (e.g. cumulants) are employed for classification. There are two critical limitations that hamper the applicability of these approaches in under-determined MIMO modrec: (i) cumulants dispersion, which deteriorates their discriminative power and (ii) inability for channel estimation, as H is under-determined.

Cumulants dispersion. Fig. 3 illustrates the effects of under-determined sensing on state-of-the-art MIMO modrec that uses cumulants [28]. We consider 200 instances of QPSK and 200 of 16-QAM, and for each instance we calculate a feature vector comprised of the seven cumulants used in prior work: $[C_{40}, C_{41}, C_{42}, C_{60}, C_{61}, C_{62}, C_{63}]$. We then adopt principle component analysis (PCA) [17] to reduce the dimensionality of instances to two dimensions and plot them in Fig. 3 for increasing determination of the sensing




Fig. 3: Discriminative power of cumulants extracted from under- (2x1), fully- (2x2) and over-determined (2x4) sensing. Red is 16-QAM, while blue, QPSK. The discriminative power of cumulants deteriorates with under-determined sensing.

system. The figure presents scatter plots of these projections for QPSK (blue) and 16-QAM (red) and the respective density “iso”-lines for each of the classes. From right to left we have over-, fully- and under-determined signal. *The increasing overlap of the classes demonstrates qualitatively that the discriminative power of cumulants deteriorates in under-determined settings.* In § 4 we quantify the effect of under-determined sensing on modrec accuracy.

Inability for channel estimation. If the channel information is not apriori known, which often occurs in a non-cooperative cognitive radio sensing, existing modrec algorithms adopt an additional step to estimate H via Independent Component Analysis (ICA) [5] or Expectation Maximization (EM) [51]. Such approaches, however, *require over-determined sensing*, and thus, impose prohibitive cost on the number of receiving antennas. Furthermore, these methods are *not directly applicable with under-determined MIMO modrec*, as they would trigger channel estimation with fewer receiving than transmitting antennas, which is unfeasible.

2.3 Under-determined MIMO sensing

To alleviate the problem of high cost, we consider under-determined MIMO sensing, whereby the number of receiver antennas is lower than that of the transmitter. Fig. 2 (left) depicts schematically the system model for under-determined MIMO sensing, where the example 2x2 MIMO link is sensed with a single-antenna sensor. We call this 2x1 sensing. Intuitively, a 2x1 sensing of a 2x2 MIMO system will gain a single observation of the MIMO constellation.

Our signal model for under-determined sensing follows the same definition as in Eq. (1): $H = [h_{ij}] \in \mathbb{C}^{N_r \times N_t}$ is the channel matrix with $i \in [1, N_r]$, $j \in [1, N_t]$ and $N_r < N_t$. For example, as illustrated in Fig. 2 (left) for the 2x1 sensing case, the H matrix is a 2x1 vector of the form $H = [h_{11} h_{12}]$, where each h_{ij} is a complex number: $h_{ij} = a_{ij} + ib_{ij}$,

A MIMO constellation is comprised of $2^{M \times N_t}$ individual symbols, where M is the modulation order (e.g. $M=2$ for QPSK, $M=3$ for 8-PSK, etc.) and N_t is the number of transmitter antennas. Since we will use the geometry of the constellation as a predictive “fingerprint” of the modulation, the spread and overlap of symbols and the tightness of a symbol’s cluster will play an important role. The spread and overlap of constellation symbols depend on the channel gain and fading, whereas the tightness of individual clusters depends on the SNR. In what follows, we explore each of these in the context of our signal model and illustrate their effects on the constellation geometry and self-similarity.




Fig. 4: 2x1 QPSK constellations at SNR 20dB (left), 10 (middle) and 5dB (right). As SNR decreases the geometry of the constellation is less-pronounced.

2.3.1 Impact of noise on the constellation geometry

The SNR of a MIMO channel affects how tightly-clustered are the $\mathbb{I}\mathbb{Q}$ samples around each constellation symbol. The lower the SNR, the more dispersed the samples and the harder it is to recognize the modulation or demodulate the signal. This effect is illustrated in Fig. 4 which depicts the constellation shape of 2x1 sensing of a QPSK signal at an SNR of 20dB (left), 10dB (middle) and 5dB (right). As the SNR decreases the geometry of the constellation is characterized with a decreasingly-pronounced clusters.

2.3.2 Impact of channel gain on constellation geometry

The gain $|h_{ij}|$ of a MIMO path (i, j) is defined as:

$$|h_{ij}| = \sqrt{a_{ij}^2 + b_{ij}^2}, \quad (2)$$

where a_{ij} and b_{ij} are the channel coefficients. We introduce the *channel gain sum* (G_S) and the *channel gain ratio* (G_R) which jointly control the MIMO constellation shape, where

$$G_S = \sum_{ij} |h_{ij}|, \quad (3)$$

controls the *spread* of the MIMO constellation symbols. High G_S allows easier MIMO decoding and modrec. The channel gain ratio G_R , defined as:

$$G_R^p = \frac{\mathcal{L}(p)}{\mathcal{L}(p+1)} = \frac{|h_{ij}^{p-1}|}{|h_{ij}^p|}, 0 \leq G_R^p \leq 1, \quad (4)$$

controls the *overlap* of constellation symbols. Here the individual paths $|h_{ij}|$ are considered in decreasing order $\mathcal{L} = [|h_{ij}^p|]$, where $p \in [1, P-1]$ and $P = N_t \times N_r$ is the total number of individual paths (i, j) .

To put the notion of G_R^p into context, let us consider our running example of 2x1 sensing from Fig. 2. Fig. 5 demonstrates the effects of G_R^p on the geometry of the constellation. In the left of the figure, we see an example of 2x1 QPSK sensing with $G_R^p = 0.1$. This means that the channel gain $|h_{12}|$ of the second link is much smaller than that of the first link. As a result the centroids of each sub-constellation are far apart (due to the high gain on the first link $|h_{11}|$), whereas the symbols in each sub-constellation are very close (due to the low gain on the second link $|h_{12}|$). Thus, the geometry of the overall constellation following 2x1 sensing of a MIMO channel with $G_R^p = 0.1$ resembles a QPSK SISO constellation. Next, we consider the constellation of 2x1 sensing with $G_R^p = 0.5$, which is illustrated in the middle pane of Fig. 5. $G_R^p = 0.5$ means that the gain of the stronger channel is twice that of the weaker (i.e. $|h_{11}| = 2 \times |h_{12}|$). For this setting, we get a “fully-unfolded” MIMO constellation with the maximum of $2^{M \times N_t}$ non-overlapping symbols, which in the case of 2x1 QPSK sensing is $2^{2 \times 2} = 16$. Finally, for $G_R^p = 1$, $|h_{11}| = |h_{12}|$, the distance between sub-constellation symbols is the same as that between the sub-constellation centroids. As a result, some of




Fig. 5: Effects of G_R^p on the MIMO constellation geometry. For $G_R^p = 0.1$ (left), where the link is dominated by one of the paths, the modulation constellation resembles SISO. For $G_R^p = 0.5$ (middle) we have a fully-unfolded MIMO constellation. For $G_R^p = 1$ (right) we have a partially-unfolded constellation. Different colors indicate meta groups.

the symbols in the sub-constellations overlap, as illustrated in the right pane of Fig. 5. This results in a partially-unfolded MIMO constellation with a total of nine non-overlapping symbols. Based on this qualitative analysis, we expect that the channel sum and ratio will affect the informativeness of geometric features for modrec; a hypothesis we confirm experimentally in §4.

2.3.3 Channel fading impact on under-determined modrec

Channel fading is a random process that models the change of the channel gain over time as a result of multipath signal propagation. In the context of modrec, of specific interest are the *channel model* and the *coherence time*, which jointly determine to what extent and how often does the channel change. Channel fading introduces time-variance in the sensed data and subsequently distorts the shape of a MIMO constellation. This, in turn, may affect the accuracy of supervised modulation classification, such as the one performed by **SYMMeTRY**.

Thus, we set forth to understand the effects of channel fading on our method (results in §4.4). Several well-established models capture the fading of wireless channels [34], including the Rayleigh model, typical for multipath environments; and the Rician model, for channels with strong Line-Of-Sight path. The channel gain $|h_{ij}| = \sqrt{a_{ij}^2 + b_{ij}^2}$ is said to be a *Rayleigh random variable*, if $a_{ij} \in \mathcal{N}(\mu_{a_{ij}}, \sigma_H)$ and $b_{ij} \in \mathcal{N}(\mu_{b_{ij}}, \sigma_H)$, such that $\mu_{a_{ij}} = \mu_{b_{ij}} = \mu_{ij}$ [34]. Similarly, $|h_{ij}|$ is said to be a *Ricean random variable* if $a_{ij} \in \mathcal{N}(\mu_{a_{ij}}, \sigma_H)$ and $b_{ij} \in \mathcal{N}(\mu_{b_{ij}}, \sigma_H)$, such that $\mu_{a_{ij}} \neq \mu_{b_{ij}}$ [34]. Our evaluation (§4) adopts these realistic models, whereby we control the channel gain by setting $\mu_{a_{ij}}$ and $\mu_{b_{ij}}$, the severity of the fading by controlling σ_H and the rate of the fading by controlling the number of consecutive samples K , for which the channel remains unchanged. We demonstrate that **SYMMeTRY** is robust to distortions introduced by realistic time-varying channels.

3 METHODOLOGY

The key insight behind **SYMMeTRY** is that the geometry of MIMO constellations exhibits nested self-similarity, which can be exploited as a classification feature. We summarize this self-similar structure into a discriminative fingerprint which can then be employed for supervised feature-based modulation recognition. We design our novel features leveraging the Minkowski-Bouligang fractal dimension [19]. We combine these features with traditional higher order cumulants [41] and evaluate their individual and joint accuracy for under-determined MIMO modrec. We demonstrate that the combined features are robust to various signal conditions and enable improvements over cumulants alone.

3.1 Preliminaries

We first present a short introduction to higher order cumulants [41] and fractal dimensions [19]. While cumulants have been previously employed in modrec, fractal analysis, which helps summarize self-similarity in IQ sample constellations, has not been considered for modrec before.

Higher Order Cumulants are a successful family of features used in the modrec literature [14], [41]. Cumulants summarize the statistical properties of IQ samples, arising from a complex-valued stationary random process $x(n)$ [41]. The k -th order cumulant is a polynomial function of moments:

$$C_{kv} = \sum_{(\cup_{p=1}^q I_p)=I} [(-1)^{q-1}(q-1)! \prod_{p=1}^q M_{\hat{k}\hat{v}}], \quad (5)$$

where $M_{kv} = E[x(n)^{k-v} x^*(n)^v]$ are the empirical estimates of the moments associated with the stationary process from which the IQ samples are drawn, and $x^*(n)$ denotes the complex conjugation of $x(n)$. The summation extends over all partitions $\{I_1, \dots, I_q\}$, $q \in \{1, \dots, k\}$. Each partition consists of q sets and \hat{k} and \hat{v} are the number of complex and conjugation terms in set p of partition I_p . Some commonly used cumulants in the literature [41] are defined as follows in terms of moments: $C_{21} = M_{21}$ and $C_{42} = M_{41} - (M_{20})^2 - 2(M_{21})^2$. In practice, fourth- and sixth-order cumulants have received most attention, and to remove the effect of the signal scale on cumulants, they are typically normalized by C_{21} [41]. In addition, since some cumulants are complex numbers, their L_2 norm is adopted as a real-valued feature in classification.

Fractal Dimensions. The intrinsic (or fractal) dimension of a finite set of points is a widely adopted descriptor to quantify fractals' self-similarity [19]. Fractal dimensions have been employed in dimensionality reduction and feature selection in machine learning [43], to optimize the utility of spatial index structures [9] and to analyze biological images [23]. While there are different kinds of dimensions that characterize fractals [36], we focus on the *Minkowski-Bouligand* (AKA *box-counting*) dimension, which was successfully employed in data analysis applications due to efficient implementations and relatively low-computational complexity [11], [21].

For a set of points in Euclidean space $X = \{x_i\}, x_i \in \mathcal{R}^d$ and a space partitioning in a grid of resolution ϵ (i.e. ϵ is the size of a hyper-cube voxel in the grid), let the *box count* $N(X, \epsilon)$ denote the number of voxels required to cover the points in X . The box counting dimension is defined as the logarithmic rate of increase of the box count $N(X, \epsilon)$ as a function of the log of the resolution $1/\epsilon$ in the limit:

$$D_{\text{box}}(X) := \lim_{\epsilon \rightarrow 0} \frac{\log N(X, \epsilon)}{\log(1/\epsilon)}. \quad (6)$$

This quantity is typically estimated numerically by varying ϵ starting from a fixed maximal resolution and estimating a linear fit for the log-log plot of $N(X, \epsilon)$ as a function of $1/\epsilon$. In this work we are not interested in the actual dimension size, but instead in the discriminative power of the cover set growth functions as features characterizing the constellations of observed IQ samples.

3.2 SYMMeTRy Overview

As other supervised classification tasks, **SYMMeTRy**, which is illustrated in Fig. 6 has two phases: training and testing.

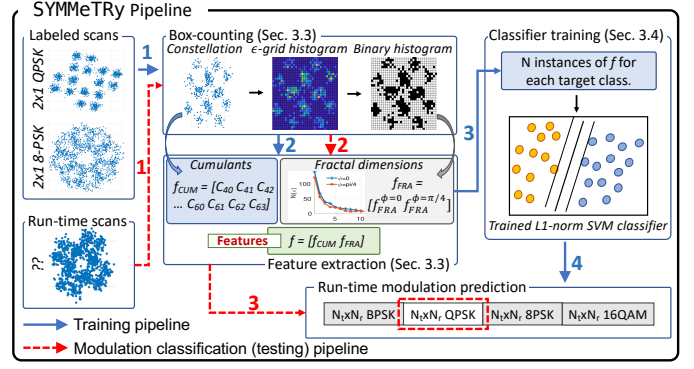


Fig. 6: Overview of **SYMMeTRy**. Blue arrows represent the training process, whereas red depict the modulation classification (i.e. testing) process.




Fig. 7: Pipeline of converting IQ samples to a binary matrix for box-counting.

Training, depicted with blue solid arrows, begins with the collection of a labeled training dataset of target modulations. Step 1 feeds the data into the *box-counting* algorithm (§ 3.3), which first converts the constellation into an ϵ -sized histogram of sample occurrences, and then thresholds it to determine which of the histogram bins are full and which empty. In step 2 (§ 3.3) we take the raw constellation and the binary histogram from step 1 and use these to extract the classification features. **SYMMeTRy** uses a combination of high order cumulants (§3.1) and fractal dimensions (§3.3). In step 3 we train a L1-norm SVM classifier (§3.4), which is then used for runtime modulation recognition.

The runtime modulation classification, illustrated with red dashed arrows, begins with the collection of an unlabeled spectrum trace. Step 1 passes the trace to the *box-counting* algorithm, which then feeds into the feature extraction in step 2. Step 3 uses the pre-trained classifier for runtime recognition of the underlying transmitter modulation. In what follows, we detail each of these steps in turn.

3.3 Extracting Fractal Features

Our goal is to employ fractal dimensions to extract signatures of modulation-specific IQ samples. Thus, we summarize the self-similar hierarchical structure of MIMO constellations using the box-counting dimension approximation. Informally, *our key idea is that samples from the same modulation will exhibit discriminative box counting growth patterns which differ from those corresponding to other modulations.*

We represent the complex IQ samples within an instance as 2-dimensional real points X based on their real and imaginary parts. Since we will consider the box counts $N(\epsilon)$ starting from a fixed ϵ_{min} resolution, we first pre-compute the number of points in each box at the highest resolution. We then aggregate the point counts at decreasing box resolution. This aggregation helps us (i) keep the complexity of subsequent steps fixed regardless of the number of IQ samples and (ii) allows us to filter noise in low-occupancy bins early in the feature extraction process.

This process is demonstrated in Fig. 7 for a 2x2 QPSK MIMO signal. We first normalize the IQ samples. We then




Fig. 8: An illustrative comparison of the fractal feature vectors for four modulations in 2-transmitter MIMO using exponential (left) and linear (right) resolution schedules T . X-axis is the box size ϵ .

impose a square grid at resolution ϵ_{min} over the IQ constellation and compute a histogram of point counts in each box (Fig. 7 middle). Finally, we binarize the histogram into full/empty (i.e. 1/0) boxes to obtain a set of coordinates containing samples (Fig. 7 right). The simplest approach to this binarization is to declare any cell with at least one point as full. In order to minimize the effect of noise, we consider a more general frequency thresholding which effectively declares “near-empty” cells as empty. We adopt Otsu’s method commonly used to binarize histograms [32]. The result of the above pre-processing is a binary matrix $Y^{k \times k}$, $y_{ij} \in \{0, 1\}$ used as an input to the box-counting.

Representing an observed IQ constellation by its estimated fractal dimension D_{box} may result in loss of discriminative power. Instead we compute the empirical decrease of $N(Y, \epsilon)$ as we increase ϵ from ϵ_{min} to a value ϵ_{max} which covers the entire Y matrix with a single box. Let $T \in \mathcal{R}^m$ be a schedule of increasing box sizes and $N(Y, T_i)$ denote the number of boxes at resolution T_i needed to cover all non-empty cells of our binarized matrix representation of the IQ samples Y . Our fractal features are comprised of the m -dimensional vector $f_{FRA} = [N(Y, T_1), \dots, N(Y, T_m)]$.

Resolution schedule T. In methods approximating the fractal dimension via box-counting, consecutive ϵ values are increased exponentially as the goal is to estimate the slope of an exponential fit of $N(X, \epsilon)$. Assuming a fixed ϵ_{min} corresponding to a size- k quantization in Y , the consecutive sizes of boxes in an exponential schedule have the following form $T = (c^i)$ for an integer $c > 1$. In order to extract a more detailed, and potentially more discriminative, shape of the $N(Y, \epsilon)$, we also consider linear schedules of the form $T = (ci)$ for integer values of $c > 1$. A comparison of the average extracted shapes using exponential $T = (2^i)$ (left) and linear $T = (i)$ (right) schedules is presented in Fig. 8. The data for this figure includes four modulations (BPSK, QPSK, 8PSK and 16QAM) assuming AWGN channel with 100 instances for each class and 300 IQ samples per instance. The x-axis presents the increasing box size ϵ , while the y-axis is the average and standard deviation of the box count across the 100 instances at a given epsilon resolution. The curves, present the box count slopes for the four modulations and illustrate (i) their discriminative power and (ii) the increased detail due to the linear schedule.

Handling phase variation. While the fractal features are invariant to changes in the amplitude within the same modulation (because we re-scale IQ samples to have a fixed maximal norm), variations in the phase, which effectively “rotate” the constellation, may result in significantly different features f_{FRA} from the same modulation. We




Fig. 9: 2x2 QPSK MIMO with a phase rotation $\phi = 0$ (left) and $\phi = \pi/4$ (right). The fractal dimensions of changing phase rotations could appear different (right). In the right figure, x-axis indicates the index of all grid size enumerations (using exponential growth) and the y-axis indicates the extracted slopes.

demonstrate this sensitivity to rotation in Fig. 9. The left and middle pane show the same constellation with a phase offset of $\pi/4$, while the right pane shows the corresponding box-counting profiles $N(Y, \epsilon)$. The reason for the differences is that covering boxes are axis-aligned, while clusters in the constellation may have irregular shapes which under rotation may require different number of boxes to be covered at the same resolution. To overcome this challenge, we rotate an observed sample at different phase offset and summarize offset-specific feature variants: f_{FRA}^ϕ .

It is important to note that noise-free (theoretical) MIMO constellations have a rotational symmetry of order 4 [45], and hence, considering phase offsets outside of $[0, \pi/4]$ would not yield additional discriminative power. Hence, it is reasonable to consider rotations in this interval, where more rotations will potentially add discriminative power at the cost of more features. Our evaluation presented in §4.11 shows that simply adding a $\pi/4$ rotation provides an accuracy boost while further rotations add negligible improvements. Thus, our fractal features in all experiments are the concatenation of these two angles $f_{FRA} = (f_{FRA}^{\phi=0}, f_{FRA}^{\phi=\pi/4})$.

3.4 Sparsity Regularized Classification

Our feature vector is constructed by concatenating cumulant and fractal features $f = [f_{CUM}, f_{FRA}]$. The features obtained may have redundant information over multiple dimensions due to multi-perspective projection during the geometric feature extraction. Therefore we utilize an 1-norm linear SVM for classification [50]. It replaces the standard ridge penalty with a lasso penalty, which enforces sparsity of coordinates for the the separation hyperplane, and thus, performs feature selection along with classification. Note that this simple classifier is not a mandatory component of our framework and can be replaced by standard feature selectors followed by any classification approach, including deep learning models.

4 EXPERIMENTAL EVALUATION

We evaluate **SYMMeTRy** in a realistic simulation and a small-scale MIMO USRP testbed. Our results show that CUM and FRA features exhibit complimentary performance. Their combination, however, maintains high and consistent performance across all noise regimes, fading conditions and with increasing MIMO transmitter complexity. These trends are retained both in simulation and with over-the-air experiments and demonstrate the high potential of **SYMMeTRy** to enable low-cost under-determined MIMO modrec.

4.1 Experimental setup

Implementation. Our box counting pipeline is implemented in MATLAB and executed on Ubuntu 14 PC. We employ liblinear’s [10] implementation of SVMs and one-vs-rest [4] training and evaluation for multi-class classification.

Data. We use data from a realistic simulation and from an over-the-air experiment in a 2x1 MIMO USRP testbed. Our synthetic datasets are generated with the MATLAB Communications System Toolbox and include four modulations: BPSK, QPSK, 8PSK and 16QAM, which are typically considered in prior modrec literature [42], [51]. Training and testing instances contain 512 IQ samples each.

Evaluation strategy. Our goal is to evaluate **SYMMeTRY** across varying constellation shapes and channel conditions. To control the constellation shapes, we vary the channel gain sum G_S (Eq. (3)) and the channel gain ratio G_R (Eq. (4)). To control the rate and severity of fading, we set the number of consecutive samples K that experience constant channel and the variance σ_H of each individual path gain $|h_{ij}| \in \mathcal{N}(\mu_{ij}, \sigma_H)$. We also consider varying SNR levels. Unless otherwise noted, we use 2x1 MIMO sensing. In §4.5 we also perform evaluation in more complex settings.

In experiments with fixed (G_S, G_R) combinations, we obtain individual channel coefficients a_{ij} and b_{ij} (Eq. (2.3)) and (2) by solving a linear system of equations (3) and (4), which for a 2x1 sensing setup with $|h_{11}| > |h_{12}|$ gives us

$$|h_{11}| = G_S * G_R / (1 + G_R) \text{ and } |h_{12}| = G_S / (1 + G_R) \quad (7)$$

Since this system is under-determined w.r.t. a_{ij} and b_{ij} , we further assume equal real and imaginary components: $a_{ij} = b_{ij}$. Thus, as per Eq. (2), $a_{ij} = |h_{ij}| / \sqrt{2}$. Finally, to evaluate **SYMMeTRY** with time-variant channels while controlling G_R and G_S , we draw the path gains $|h_{ij}|$ from a normal distribution $\mathcal{N}(\overline{h_{ij}}, \sigma_H)$. For a 2x1 setup, we set $\overline{h_{11}}$ and $\overline{h_{12}}$ according to Eq.(7) and vary σ_H .

In all experiments we compare the accuracy of **SYMMeTRY**, defined as the fraction of correctly-predicted instances over all instances. We compare the accuracy across three different feature configurations: cumulants (CUM) which are employed in all classification-based prior modrec work, fractals (FRA) and their combination (CUM+FRA).

4.2 SYMMeTRY on a time-invariant channel

We first evaluate the performance of **SYMMeTRY** in time-invariant channels, i.e. H is fixed in all instances. Our results indicate that CUM and FRA alone have complementary advantages in different gain/noise regimes. The combination of the two features maintains consistently advantageous performance across all experimental regimes.

4.2.1 Effects of channel gain ratio G_R

We first evaluate the effects of symbol overlap controlled by G_R . As detailed in §2, with a small G_R the MIMO constellation is folded, converging to its SISO equivalent. Mid-range G_R produces a fully-unfolded MIMO constellation, while G_R close to 1 produces a partially-unfolded constellation, whereby some constellation symbols overlap. These effects of G_R on the constellation geometry directly affect the discriminative power of our features. To quantify this, we fix the channel noise $w(n)$ (Eq. (1)) and scale G_R from 0 to 1 in increments of 0.1 in a 2x1 sensing setup. Figs. 10a, 10b present our results comparing the accuracy in mid-noise (10dB) and high noise (5dB) settings, respectively. The performance of cumulants deteriorates as G_R increases. This is expected, as cumulants are robust with SISO modrec (i.e. when G_R is low) and deteriorate as the MIMO constellations




Fig. 10: Accuracy over varying channel gain ratio G_R with time-invariant channel at medium SNR of 10dBm (a) and low SNR of 5dBm (b). Cumulants are robust in SISO-like cases (i.e. at low G_R , whereas fractals perform better in MIMO-like cases (mid to high G_R). Thus, the combination of the two features leads to robust performance classification across constellation shapes

unfold with increasing G_R . The fractal features retain stable and high performance across all G_R settings for SNR=10dB, which demonstrate the robustness of the feature to channel variations. With lower SNR of 5dB, cumulants outperform fractals at low G_R and switch at high G_R . The CUM+FRA combination retains a stable and high performance across all ratios even when the SNR is 5dB. These results indicate the robustness of the combined feature CUM+FRA to noise and constellation shape changes.

4.2.2 Effects of channel gain sum G_S

The sum of individual path gains affects the spread of the MIMO constellation: the higher the gain, the more spread the symbols are, the easier will be to classify a signal's modulation. Thus, we set out to evaluate the effects of G_S on **SYMMeTRY**'s performance. We vary G_S from 0.2 to 2 in increments of 0.2 for 2x1 sensing. Figs. 11a, 11b present our results for mid-noise (10dB) and high noise (5dB) regimes. Across all regimes, FRA and the combined CUM+FRA features outperform cumulants alone. For low gain regimes all counterparts suffer deteriorated performance, which rebounds as the sum gain increases beyond 1. These trends are consistent across SNR regimes.




Fig. 11: Accuracy over varying channel gain sum G_S with time-invariant channel at medium SNR of 10dBm (a) and low SNR of 5dBm (b). FRA and CUM+FRA features outperform CUM across all G_S . Performance is consistent across SNR regimes.

4.2.3 Effects of SNR

We now seek to evaluate the effects of channel noise on **SYMMeTRY**'s performance. We vary the SNR from 0 to 20dB in increments of 5dB for two channels $H1$ and $H2$ ($G_S^{H1} > G_S^{H2}$). The channels were generated with the MATLAB `comm.MIMOchannel` block and their respective channel responses were $H1 = [1.17 - 0.32i, 1.03 + 0.36i]$ and $H2 = [0.74 + 0.17i, 0.75 - 0.13i]$. Our results are presented in Figs. 12a 12b. FRA and FRA+CUM outperform CUM across all SNR regimes for both channel conditions. With high-gain channels, (i.e. $H1$ in Fig. 12a), the performance of FRA and CUM+FRA over CUM is further emphasized.

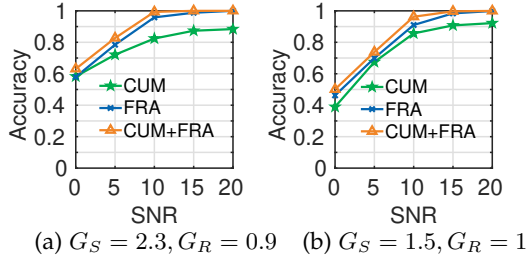


Fig. 12: Accuracy with varying SNR for two channels $H1$ (a) and $H2$ (b). The channel gain sum G_S for $H1$ is larger than that for $H2$ (i.e. constellations from $H1$ are spread out and easier to detect than $H2$). FRA and FRA+CUM outperform CUM across all SNR regimes. High-gain channels ($H1$), further emphasize the advantage of FRA and CUM+FRA over classical CUM.

4.3 SYMMeTRY on a time-varying channel

In reality, channels vary with time depending on changes in a link's environment. Hence, we next evaluate the effects of time-varying channels on **SYMMeTRY**'s performance. In our model (§2.3.3), we control fading by setting the variance σ_H of the path gain, and the speed at which the channel varies, by setting the number of consecutive samples K for which the channel is time-invariant. We begin by evaluating the effects of σ_H and K on **SYMMeTRY**'s performance. We then evaluate performance with changing constellation geometry (i.e. G_R and G_S). Finally, we examine the effects of SNR. Our findings show that the performance of both FRA and CUM alone deteriorates with a time-varying channel, however, the combined feature maintains high performance across varying constellation shapes, fading rates and noise level.

4.3.1 Effects of channel fading

We now set G_R to 0.8, G_S to 3 and K to 1 and evaluate accuracy while increasing the path gain variance σ_H from 0.1 to 1 in increments of 0.1. Figs. 13a, 13b show our results for SNR 10dB(left) and 5dB(right). FRA outperforms CUM at lower channel variance, whereas the contrary holds with high variance. The combined feature consistently outperforms either of the features alone in both SNR regimes.




Fig. 13: Accuracy with increasing channel fading as controlled by σ_H for mid-SNR 10dBm (a) and a low-SNR 5dBm (b) settings. FRA+CUM consistently outperforms CUM/FRA alone across all fading and SNR regimes.

4.3.2 Effects of channel coherence time

The channel coherence time captures the duration for which the channel response remains unchanged. In our simulation, we control this by setting K , which determines the number of consecutive instances during which the channel remains constant. We define an instance as a vector of 512 IQ samples. Figs. 14a, 14b present **SYMMeTRY**'s accuracy at SNR 10 and 5dB as K increases from 1 to 11 in increments of 1. G_R is 0.8, G_S is 1.5 and σ_H is 0.1. The combined feature FRA+CUM outperforms the other two features consistently in both SNR regimes. The performance does not change significantly, as K increases, indicating that **SYMMeTRY** will perform robustly in both fast- and slow-fading channels.

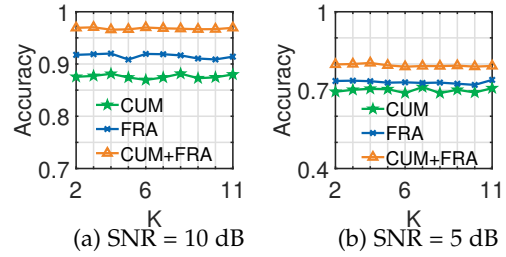


Fig. 14: Accuracy with increasing channel coherence time for a mid-SNR 10dBm (a) and low-SNR 5dBm (b) setting. Performance is consistent across K and SNR, indicating that **SYMMeTRY** is robust in both fast- and slow-fading channels.

4.3.3 Effects of the constellation geometry

We now evaluate the effects of the constellation shape. As before, we control the overlap of constellation symbols with G_R and the spread of the constellation with G_S . Figs. 15a, 15b present our results for varying G_R with $G_S = 3$, $K = 1$ and $\sigma_H = 0.1$. FRA outperforms CUM at SNR=10dB, while CUM outperforms FRA for G_R lower than 0.5 at SNR=5dB. The combined feature leads to high and consistent performance across both SNR regimes and all G_R settings. In addition, Figs. 15c, 15d evaluate the effects of constellation spread (as controlled by G_S) on performance. We vary G_S from 0.2 to 2 in steps of 0.2. G_R is 0.8, K is 1 and σ_H is 0.1. The combined CUM+FRA feature marginally outperforms the other two. The accuracy is affected by both the channel gain sum and the SNR, however, for a gain higher than 1 across both SNR regimes **SYMMeTRY** achieves high classification performance.

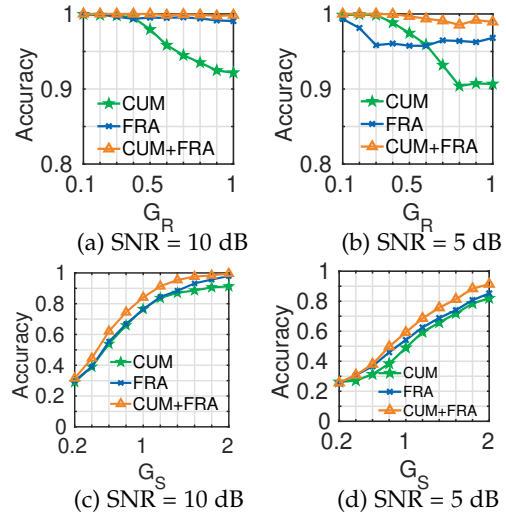


Fig. 15: Effects of constellation geometry as controlled by the channel gain ratio G_R (controlling symbol overlap) and the channel gain sum G_S (controlling symbol spread) in time-variant channels. (a), (b) show accuracy as a function of G_R for mid- and low-SNR. (c), (d) show accuracy as a function of G_S for mid- and low-SNR. FRA+CUM outperforms counterparts across all settings.

4.4 SYMMeTRY with realistic channel conditions

In this section, we remove any constraints on the constellation geometry and fading (i.e. all of G_R , G_S , σ_H and K are uncontrolled) and evaluate the performance of **SYMMeTRY**. We adopt two realistic and commonly-used channel models: Rician and Rayleigh, and evaluate the performance as a function of the channel SNR. We use the MATLAB `comm.MIMOchannel` toolbox for our implementation with the default parameter setting of the function. In brief, we assume a flat fading channel and also do not consider the

Doppler effect (i.e. the sensor and transmitter were considered stationary). We assume the path delay is zero and the averaged power of the path gains is normalized to 0 dB across realizations. We generate one dataset using Rayleigh fading distribution model and another one using Rician model. The channel for each instance is independently generated using the aforementioned MATLAB function. Figs. 16a, 16b show the results of this analysis. CUM and FRA alone have distinct regimes of high performance across both channel models. We hypothesize this is due to the complimentary nature of the two features. Inherently, FRA captures the shape of the constellation (symbols' location and empty spaces) with underdetermined sensing, however it faces challenges due to constellation dispersion with channel effects. CUM, on the other hand, are less discriminative with underdetermined sensing, however, they are more robust to real channel conditions. Thus, the combination of the two brings the best of both worlds: capturing the constellation shape with underdetermined sensing, while being robust to channel conditions.

So far, our results considered same-channel training and testing. We now explore the effects of cross-channel training on **SYMMeTRy**'s performance, focusing on four cases (Fig. 16c): (i) training on Rayleigh and testing on Rician (green), (ii) training on Rician and testing on Rayleigh (blue), (iii) training on an equal mix of Rayleigh and Rician and testing on Rician (orange) and (iv) training on the aforementioned mix and testing on Rayleigh (red). **SYMMeTRy** with cross-channel training results in high and consistent performance, commensurate with that of same-channel training. Modulation classification under Rayleigh conditions remains slightly more challenging than Rician conditions. Finally, the method is robust to both single- and mix-channel training, which demonstrates that **SYMMeTRy** is not sensitive to changing channel conditions from training to testing, and thus, the training process does not need to capture diverse channel conditions.




Fig. 16: Accuracy with increasing SNR over a Rayleigh (a) and Rician (b) channel with 2x1 sensing. **SYMMeTRy** with the combined CUM+FRA features outperforms counterparts across all SNR regimes in both channel conditions. (c) Accuracy of the CUM+FRA feature with cross-channel training. **SYMMeTRy**'s performance remains high and consistent with the same-channel training.

4.5 SYMMeTRy with increasing transmitter antennas

We now set out to examine **SYMMeTRy**'s performance with increasing transmitter complexity. For this experiment, we use the same data generation setup as discussed in the previous section (§4.4). We increase the number of transmitter antennas N_t to 3 and 4, while maintaining a single antenna on the sensor. Figs. 17a-17d demonstrate our results. Across all SNR levels and both channel models, the combined feature CUM+FRA maintains the best performance. The performance improves in high SNR regimes. Finally, even though the complexity of the transmitter affects the modrec accuracy, **SYMMeTRy** with CUM+FRA feature is able to

achieve an accuracy of over 0.8 at SNR 10dB and 0.9 at SNR 20dB for the most challenging case of 4x1 MIMO sensing. This demonstrates **SYMMeTRy**'s performance to support modrec of complex MIMO systems even when the sensor is only equipped with a single antenna.

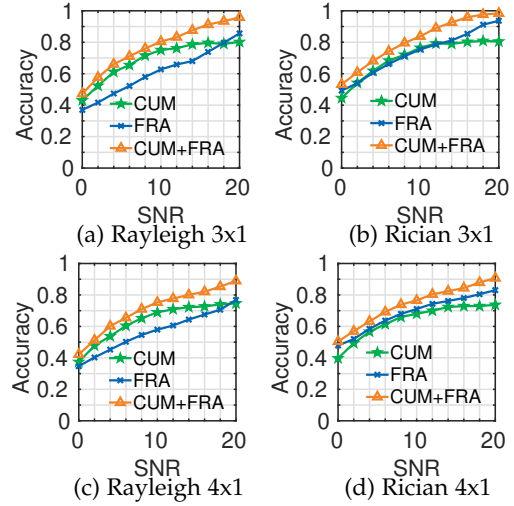


Fig. 17: Accuracy across SNR and channel models with increasingly-complex MIMO transmitter. (a) and (b) present results for 3x1 sensing, while (c) and (d) for 4x1 sensing. The combined CUM+FRA feature is able to retain high performance even when a 4-antenna transmitter is sensed with a single-antenna sensor.

4.6 SYMMeTRy with diversity coding MIMO

As discussed in the background section §2.1, depending on how data is mapped across transmitter streams, a MIMO transmitter can be realized with spatial diversity or spatial multiplexing. So far, all our experiments consider MIMO with spatial multiplexing, whereby a unique portion of the data stream is transmitted through each antenna. In this section, we evaluate the efficiency of **SYMMeTRy** with spatial diversity coding whose purpose is to exploit data redundancy across transmitter streams to improve the reliability of data transfer. We use Alamouti's code to implement MIMO transmitters with spatial diversity coding. Fig. 18 presents our results across SNR regimes for two realistic channel settings: Rayleigh (18a) and Rician (18b). **SYMMeTRy** with combined CUM+FRA features consistently outperforms counterparts across all SNR regimes and with both realistic channel settings. Furthermore, the performance is very similar to that for a MIMO system with spatial multiplexing (Fig. 16).

4.7 SYMMeTRy with MIMO precoding

As detailed in §2.1, MIMO can be either direct-mapped or precoded, depending on how transmit power and modulation are assigned to the individual transmitter streams. So far, our evaluation focused on direct-mapped MIMO, whereby each antenna stream uses the same transmit power and modulation. In this section, we evaluate **SYMMeTRy**'s performance with precoded MIMO; that is each stream is assigned a different modulation (power remains fixed across streams). We explore three combinations for a 1x2 MIMO sensing setup across SNR: BPSK+QPSK, QPSK+8PSK and 8PSK+16QAM, where training and testing are performed on the same combination (i.e. training is combination-aware). Fig. 19 presents our results for a Rayleigh and Rician fading channels. **SYMMeTRy**'s performance with the combined




Fig. 18: Accuracy with increasing SNR over a Rayleigh (a) and Rician (b) channel with 2x1 sensing. The MIMO transmitter uses spatial diversity coding with Alamouti's code. **SYMMeTRy** with the combined CUM+FRA features outperforms counterparts across all SNR regimes in both channel conditions.




Fig. 19: Accuracy for MIMO with precoding with increasing SNR over a Rayleigh (a) and Rician (b) channel with 2x1 sensing. We explore three realizations of a 2x2 MIMO transmitter with precoding: BPSK+QPSK, QPSK+8PSK and 8PSK+16QAM. **SYMMeTRy** with the combined CUM+FRA features outperforms counterparts across all SNR regimes in both channel conditions.

CUM+FRA feature remains high, albeit slightly deteriorated in comparison with direct-mapped MIMO (Fig. 16a, 16b).

While these results demonstrate **SYMMeTRy**'s power in classifying just any possible MIMO realization, it is important to note that with MIMO precoding, the training space explodes combinatorially with the transmitter complexity (i.e. the number of transmitter antennas) and the number of target modulation classes. This may pose prohibitive overhead in training data collection, which can be tackled through alternative classification frameworks (e.g. transfer learning). We leave such an exploration for future work.

4.8 SYMMeTRy with increasing modulation complexity

All experiments so far employ BPSK, QPSK, 8PSK and 16QAM with increasing transmitter complexity. We now set forth to evaluate **SYMMeTRy**'s performance with complex modulations. To this end, we consider three higher-order modulations: 32-, 64- and 128QAM. Fig. 20 presents average accuracy across the three classes for 2x1 under-determined sensing. Our results show a deterioration in the average classification performance in comparison with simpler modulations. A closer look at the confusion matrices (not shown in interest of space) indicates that **SYMMeTRy** confuses 32 and 128QAM, whereas 64QAM retains high performance. Classifying higher order modulations is challenging as their constellations are extremely dense (e.g. a 2x2 MIMO setup using 128QAM can result in up to 512 constellation points depending on channel conditions). This explains, in part, the slightly lower performance of the combined CUM+FRA feature compared to the CUM-only counterpart, as fractals are likely becoming sensitive to this extreme constellation density. We note that further exploration of **SYMMeTRy**'s parameters (e.g. grid resolution) with high-order modulations could gain improved performance. We leave such exploration for future work.




Fig. 20: Accuracy with increasing SNR over a Rayleigh (a) and Rician (b) channel for high-order modulations (32/64/128QAM).




Fig. 21: Accuracy across SNR and channel models with increasingly-complex MIMO receivers.

4.9 SYMMeTRy in determined/over-determined sensing

We now set forth to evaluate the performance of **SYMMeTRy** in determined and over-determined sensing scenarios. We seek to understand whether additional sensor antennas add any performance gain in comparison with single-antenna sensing. Similar to the under-determined experiment setup in 4.4, we evaluate the performance in MATLAB simulation with realistic channel models. We use feature concatenation to combine features extracted from scans from individual receiver streams. Notice that although the determined setting would allow it, we did not do signal recovery/channel estimation pre-processing. Instead we use the raw signal on the receiver as the input, as we want to retain the low complexity of the framework pipeline.

Figure 21 shows our results. The figure presents classification accuracy over increasing SNR for two realistic channel conditions (Rayleigh and Rician) with determined (Fig. 21a and 21b) and over-determined (Fig. 21c and 21d) sensing. CUM+FRA and FRA are two variants of **SYMMeTRy** using the combined cumulant and fractal feature and fractals only, respectively. As with the prior evaluation, CUM is a counterpart from the literature, which uses cumulants only as a feature. For the determined 2x2 setting, the combined CUM+FRA feature still holds a noticeable advantage over the CUM one across all SNR regimes and both channel models. As the number of sensor antennas grows to four (over-determined sensing), the CUM feature and **SYMMeTRy** with the CUM+FRA feature have equivalent performance. In all settings, **SYMMeTRy** with the FRA feature alone is significantly disadvantaged in determined and over-determined settings compared to CUM+FRA and CUM.

	Rx Gain	CUM	FRA	CUM+FRA
20 in.	40 dB	0.51	0.55	0.62
	50 dB	0.62	0.74	0.80
60 in.	40 dB	0.51	0.67	0.73
	50 dB	0.48	0.71	0.77

TABLE 2: Evaluation in a 2x1 MIMO USRP testbed

Considering all results so far (§4.2-§4.9), we note that our novel fractal feature gains significant advantage over existing cumulant-based counterparts in under-determined MIMO sensing, even when the number of transmitter antennas is four times greater than these at the sensor. At the same time, it does not lead to deteriorated classification performance in fully-determined and over-determined sensing scenarios considered in prior work. These results show that **SYMMeTRY** with the combined FRA+CUM feature can be used in a universal MIMO modulation classification framework towards cost-efficient and agile recognition with mixed-complexity spectrum sensors.

4.10 SYMMeTRY in real-world under-determined sensing

We evaluate the performance of our method in real over-the-air transmissions from a USRP-based testbed. We use a transmitter comprised of a USRP B210 attached to an Intel i7-5600U CPU host², and a receiver comprised of a USRP B210 with an Intel i7-6700 CPU host. Both hosts are running on low-latency Linux kernel. We establish a 2x1 sensing setup by transmitting on both Tx channels and receiving on only on Rx channel. Using GNURadio, the transmitter generates a MIMO signal modulated with BPSK, QPSK, 8PSK and 16QAM. We record 3000 samples for each modulation. We use one third of the samples as the testing set and the rest for classifier training. The two USRP devices are located in line-of-sight. We evaluate two different scenarios: one where the transmitter/sensor are 20 inches apart and another where they are 60 inches apart. For each distance, we collect two traces while setting the sensor’s Rx gain to 40 and 50dB.

Table 2 shows the results. For all settings, FRA outperform CUM by a margin of 0.04 (at 20in/40dB) to 0.23 (at 60in/50dB). The combined features maintain the maximum performance across all measurement scenarios with the highest of 0.797 at 20in and 50dB gain.

4.11 Effects of parameter selection

We now evaluate the impact of input parameters on performance. Three parameters are worth noting: the finest grid resolution M for box-counting, the number of angles ϕ used in handling phase variations (§3.3) and the sparsity regularization parameter C of the SVM classifier (§3.4).

Grid resolution for fractal dimension extraction. Fig. 22(left) present accuracy results for **SYMMeTRY** using the FRA feature across changing M . We use three settings for M ($M = \{32, 64, 128\}$). As the figures show, **SYMMeTRY**’s performance is not influenced by the selection of M .

Number of angle projections. As mentioned before, we use multiple angle projections to construct the fractal feature

². Note that USRP B210 can support up to 2-transmitter MIMO with its internal capabilities. Further increase of the transmitter complexity requires additional B210 boards and an external GPSDO for time sync.




Fig. 22: Effects of parameter setting on **SYMMeTRY**’s performance. (left) Impact of initial grid granularity (i.e. number of grid cells on each edge) on accuracy. (middle) Impact of the number of angle rotations. (right) Impact of the SVM regularization parameter setting for 2x1 sensing with increasing SNR and Rician fading.

in order to make the feature model robust to unknown rotations of the IQ constellation. We evaluate the impact on performance with increasing number of angle projections. We start from 1 angle (0 degrees) and gradually add $\frac{\pi}{2}$, $\frac{\pi}{3}$, $\frac{\pi}{4}$, $\frac{\pi}{5}$ degrees and so on. We then observe the classification accuracy of **SYMMeTRY** using the FRA feature. Figure 22 (middle) presents classification accuracy with increasing number of angle projections in three SNR regimes. For all SNR regimes, increasing the number of angle projections beyond 3 does not offer additional performance improvement.

Fine-tuning the SVM regularization parameter C . Learning an SVM model, such as the one used by **SYMMeTRY**, involves a trade-off between training error and the margin width of the classifier which translates to generalization and lower testing error. The regularization parameter C controls this trade-off and has to be set prior to classifier training. Larger values of C penalize the loss function (i.e. the testing error), whereas smaller values of C penalize the regularization function (i.e. the training error). In this experiment, we explore the effect of C setting on **SYMMeTRY**’s classification performance. To this end, we compare the accuracy for $C = 1$ and that of an optimal C selected by grid-search over an exponentially-increasing set of C values from 0.01 to 100. Fig. 22 (right) presents classification accuracy comparison of **SYMMeTRY** with the FRA features over increasing SNR with a realistic Rician channel. The green line presents results for fixed C , whereas the blue presents results for optimal grid-search-based C at a given SNR. The results show that both fixed and optimal C parameter settings lead to similar performance, indicating that **SYMMeTRY** is not sensitive to the setting in different SNR regimes. As a result we perform all remaining experiments using a fixed $C = 1$.

4.12 Runtime evaluation

In this section we evaluate **SYMMeTRY**’s runtime overhead focusing on feature extraction and classifier training as a function of the instance size. For each of our target modulations (BPSK, QPSK, 8PSK and 16QAM), we generate an increasing number of instances from 1 to 1024 on an exponential schedule. Each instance consists of 512 IQ samples. Fig. 23 presents the aggregate runtime across the four modulations for feature extraction (left) and SVM classifier training (right). The runtimes were obtained with our single-thread MATLAB implementation on a Linux machine with a 2GHz Intel Xeon processor. The feature extraction time is in the order of seconds, whereas the classifier training is in the order of milliseconds. Both runtimes grow linearly as the number of instances increases. The maximum incurred




Fig. 23: **SYMMeTRy** runtime for feature extraction (left) and classifier training (right). Both feature extraction and classifier training time grow exponentially with the instance size.

time for feature extraction and training combined is 88.092s with 1024 instances. Our evaluation employs 1000 instances. Although we have not evaluated this quantitatively, we believe that the number of input instances can be further reduced without negative implications on the classification performance, which will speed up the execution. These results illustrate **SYMMeTRy**'s potential to support near-real-time modulation recognition.

5 RELATED WORK

Modrec without channel estimation. Most previous MIMO modrec work requires channel estimation [5], [16], [18], [27], [30], and hence the constraint of high number of receivers to ensure that the linear system at the core of channel estimation techniques is not under-determined. [14] stand out from the above, as they evaluated the utility of high order statistics as features without channel estimation and established that these features are not robust to noise and the channel mixing effect. Tian et Al. [42] employed a shape-constrained clustering approach within a likelihood based framework. This work, however, incurs a very high computational cost which further grows with the constellation order. In addition, different from us, all above methods require over-determined sensing, and thus, high overhead on the sensing hardware.

Fractal Geometry. Fractal dimension analysis, which is at the core of our novel features, has been successfully employed in digital image processing with applications to medical image analysis. [15] extracts fractal dimension features from pathological images and employs it for cancer classification. In [20], an automatic scar quantification approach based on segmentation-based fractal texture analysis has been presented that provides accurate and consistent results for MRI scan sequences. Different from the above, our fractal features framework goes beyond the fractal dimension (D_{box}) and employs the detailed rate of decrease of the box counts $N(\epsilon)$ as a modulation classification feature.

Sparsity and classification. Our fractal features are in a sense exhaustive and in order to focus on the most discriminative ones we employ sparsity-promoting classifiers. The merits of sparsity on features selection have been studied extensively [46]. We adopt an L1-norm regularized linear SVM [49]. Other modrec approaches have also recently employed feature selection for exhaustive features such as order statistics [47]. The recent success of deep learning techniques across application domains have also inspired approaches for modrec in the communications domain [25], [31], [35]. Our work is complementary with advanced classification approaches as it introduces geometry-aware discriminative features which can be employed in any classification scheme, including deep learning approaches.

6 DISCUSSION AND CONCLUSION

In this paper we were the first to consider the challenging problem of supervised MIMO modulation recognition for *under-determined* spectrum sensing infrastructures. Our key idea is to exploit the inherent multi-scale self-similarity of MIMO modulation IQ constellations, which persists in under-determined settings. We set out with a careful characterization of the various effects of the channel on the the constellation organization which informed our feature-extraction framework exploiting self-similarity patterns in MIMO constellations. Our framework is rooted in the rich methodology of fractal geometry. We performed an extensive evaluation in a realistic simulation and in a USRP testbed, and demonstrated high and persistent performance across various SNR regimes, channel fading conditions and with increasing complexity of the MIMO transmitter. These results were also confirmed through under-determined MIMO sensing in a small-scale USRP testbed. Our analysis focuses on under-determined modrec of a direct-mapped MIMO link, i.e., a link in which each stream is transmitted with the same modulation. We expect that **SYMMeTRy** will have similar advantages in precoded MIMO links as well; a direction we plan to explore in the future. Our exploration and analysis paints a clear and feasible path to practical and cost-efficient sensing with few receiver antennas—an essential component of future Dynamic Spectrum Access technology, policy and enforcement.


7 ACKNOWLEDGEMENTS

This work was supported by NSF CAREER grant CNS-1845858 and NSF Smart and Connected Communities (SC&C) grant CMMI-1831547.


REFERENCES

- [1] Microsoft's Spectrum Observatory. <https://observatory.microsoftspectrum.com/>.
- [2] H. Abuella and M. K. Ozdemir. Automatic modulation classification based on kernel density estimation. *Canadian Journal of Electrical and Computer Engineering*, 39(3):203–209, 2016.
- [3] G. Atia, A. Sahai, and V. Saligrama. Spectrum enforcement and liability assignment in cognitive radio systems. In *2008 3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, pages 1–12. IEEE, 2008.
- [4] C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [5] V. Choqueuse, S. Azou, K. Yao, L. Collin, and G. Burel. Blind modulation recognition for mimo systems. *MTA Review*, 2009.
- [6] M. Di Renzo, H. Haas, A. Ghayeb, S. Sugiura, and L. Hanzo. Spatial modulation for generalized mimo: Challenges, opportunities, and implementation. *Proceedings of the IEEE*, 102(1):56–103, 2013.
- [7] O. A. Dobre, A. Abdi, Y. Bar-Ness, and W. Su. Survey of automatic modulation classification techniques: classical approaches and new trends. *IET communications*, 1(2):137–156, 2007.
- [8] O. A. Dobre, Y. Bar-Ness, and W. Su. Higher-order cyclic cumulants for high order modulation classification. In *IEEE MILCOM*, Boston, MA, 2003.
- [9] C. Faloutsos and I. Kamel. Beyond uniformity and independence: Analysis of r-trees using the concept of fractal dimension. In *Proceedings of the thirteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 4–13. ACM, 1994.
- [10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *Journal of machine learning research*, 9(Aug):1871–1874, 2008.
- [11] K. Foroutanpour, P. Dutilleul, and D. L. Smith. Advances in the implementation of the box-counting method of fractal dimension estimation. *Applied mathematics and computation*, 1999.

- [12] H. Gang, L. Jiandong, and L. Donghua. Study of modulation recognition based on HOCs and SVM. In *IEEE VTC Spring*, Milan, Italy, 2004.
- [13] L. Han, F. Gao, Z. Li, and O. A. Dobre. Low complexity automatic modulation classification based on order-statistics. *IEEE Trans. on Wireless Communications*, 16(1):400–411, 2017.
- [14] K. Hassan, I. Dayoub, W. Hamouda, C. N. Nzeza, and M. Berbineau. Blind digital modulation identification for spatially-correlated mimo systems. *IEEE Trans. on Wireless Communications*, 11(2):683–693, 2012.
- [15] P.-W. Huang and C.-H. Lee. Automatic classification for pathological prostate images based on fractal analysis. *IEEE Trans. on Medical Imaging*, 28(7):1037–1050, 2009.
- [16] S. Huang, Y. Yao, Z. Wei, Z. Feng, and P. Zhang. Automatic modulation classification of overlapped sources using multiple cumulants. *IEEE Trans. on Vehicular Technology*, 66(7), 2017.
- [17] I. Jolliffe. *Principal component analysis*. Springer, 2011.
- [18] E. Kanterakis and W. Su. Modulation classification in mimo systems. In *Proc. IEEE MILCOM*, pages 35–39, 2013.
- [19] K. Kumaraswamy. Fractal dimension for data mining. *Center for Automated Learning and Discovery School of Computer Science Carnegie Mellon University*, 5000, 2003.
- [20] T. Kurzdorfer, K. Breininger, S. Steidl, A. Brost, C. Forman, and A. Maier. Myocardial scar segmentation in lge-mri using fractal analysis and random forest classification. In *IEEE 24th International Conference on Pattern Recognition (ICPR)*, 2018.
- [21] L. S. Liebovitch and T. Toth. A fast algorithm to determine fractal dimensions by box counting. *Physics Letters A*, 141(8-9), 1989.
- [22] G. Lu, K. Zhang, S. Huang, Y. Zhang, and Z. Feng. Modulation recognition for incomplete signals through dictionary learning. In *IEEE WCNC*, pages 1–6, San Francisco, CA, 2017.
- [23] W. Lu, Y. Tong, Y. Yu, Y. Xing, C. Chen, and Y. Shen. Applications of artificial intelligence in ophthalmology: General overview. *Journal of ophthalmology*, 2018, 2018.
- [24] M. A. McHenry, P. A. Tenhula, D. McCloskey, D. A. Roberson, and C. S. Hood. Chicago Spectrum Occupancy Measurements and Analysis and a Long-term Studies Proposal. In *ACM TAPAS*, Boston, MA, 2006.
- [25] F. Meng, P. Chen, L. Wu, and X. Wang. Automatic modulation classification: A deep learning enabled approach. *IEEE Trans. on Vehicular Technology*, 67(11):10760–10772, 2018.
- [26] R. Y. Mesleh, H. Haas, S. Sinanovic, C. W. Ahn, and S. Yun. Spatial modulation. *IEEE Transactions on vehicular technology*, 57(4), 2008.
- [27] M. S. Muhlhaus, M. Oner, O. A. Dobre, and F. K. Jondral. A low complexity modulation classification algorithm for mimo systems. *IEEE Communications Letters*, 17(10):1881–1884, 2013.
- [28] M. S. Muhlhaus, M. Öner, O. A. Dobre, H. U. Jkel, and F. K. Jondral. Automatic modulation classification for mimo systems using fourth-order cumulants. In *2012 IEEE Vehicular Technology Conference (VTC Fall)*, pages 1–5. IEEE, 2012.
- [29] A. K. Nandi and E. E. Azzouz. Modulation recognition using artificial neural networks. *Signal processing*, 56(2):165–175, 1997.
- [30] Y. Nie, X. Shen, S. Huang, Y. Zhang, and Z. Feng. Automatic modulation classification based multiple cumulants and quasi-newton method for mimo system. In *Proc. IEEE WCNC*, 2017.
- [31] T. J. O’Shea, T. Roy, and T. C. Clancy. Over-the-air deep learning based radio signal classification. *IEEE Journal of Selected Topics in Signal Processing*, 12(1):168–179, 2018.
- [32] N. Otsu. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1):62–66, 1979.
- [33] P. Panagiotou, A. Anastasopoulos, and A. Polydoros. Likelihood ratio tests for modulation classification. In *IEEE MILCOM*, Los Angeles, CA, 2000.
- [34] J. G. Proakis and D. G. Manolakis. *Digital signal processing 4th edition*, 2006.
- [35] S. Rajendran, W. Meert, D. Giustiniano, V. Lenders, and S. Pollin. Deep learning models for wireless signal classification with distributed low-cost spectrum sensors. *IEEE Trans. on Cognitive Communications and Networking*, 2018.
- [36] S. N. Rasband. *Chaotic dynamics of nonlinear systems*. Courier Dover Publications, 2015.
- [37] S. Roy, K. Shin, A. Ashok, M. McHenry, G. Vigil, S. Kannam, and D. Aragon. Cityscape: A metro-area spectrum observatory. In *IEEE ICCCN*, Vancouver, B.C., Canada, 2017.
- [38] V. Sathya, S. M. Kala, M. I. Rochman, M. Ghosh, and S. Roy. Standardization advances for cellular and wi-fi coexistence in the unlicensed 5 and 6 ghz bands. *GetMobile: Mobile Computing and Communications*, 24(1):5–15, 2020.
- [39] M. M. Sohel, M. Yao, T. Yang, and J. H. Reed. Spectrum access system for the citizen broadband radio service. *IEEE Communications Magazine*, 53(7):18–25, 2015.
- [40] N. Soltani, K. Sankhe, S. Ioannidis, D. Jaisinghani, and K. Chowdhury. Spectrum awareness at the edge: Modulation classification using smartphones. In *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 1–10. IEEE, 2019.
- [41] A. Swami and B. M. Sadler. Hierarchical digital modulation classification using cumulants. *IEEE Trans. on Communications*, 48(3):416–429, 2000.
- [42] J. Tian, Y. Pei, Y.-D. Huang, and Y.-C. Liang. Modulation-constrained clustering approach to blind modulation classification for mimo systems. *IEEE Trans. on Cognitive Communications and Networking*, 4(4):894–907, 2018.
- [43] C. Traina Jr, A. Traina, L. Wu, and C. Faloutsos. Fast feature selection using fractal dimension. *Journal of Information and data Management*, 1(1):3, 2010.
- [44] B. Van den Bergh, D. Giustiniano, H. Cordobés, M. Fuchs, R. Calvo-Palomino, S. Pollin, S. Rajendran, and V. Lenders. Electrosense: Crowdsourcing spectrum monitoring. In *Proc. IEEE DySPAN*, pages 1–2, 2017.
- [45] N. Wang, W. Liu, H. Men, M. Jin, and H. Xu. Further complexity reduction using rotational symmetry for edas in spatial modulation. *IEEE Communications Letters*, 18(10):1835–1838, 2014.
- [46] L. Wolf and A. Shashua. Feature selection for unsupervised and supervised inference: The emergence of sparsity in a weight-based approach. *Journal of Machine Learning Research*, 6(Nov), 2005.
- [47] Z. Wu, S. Zhou, Z. Yin, B. Ma, and Z. Yang. Robust automatic modulation classification under varying noise conditions. *IEEE Access*, 5:19733–19741, 2017.
- [48] W. Xiong, P. Bogdanov, and M. Zheleva. Robust and efficient modulation recognition based on local sequential iq features. In *IEEE International Conference on Computer Communications (IEEE INFOCOM 2019)*, 2019.
- [49] L. Zhang and W. Zhou. On the sparseness of 1-norm support vector machines. *Neural Networks*, 23(3):373–385, 2010.
- [50] J. Zhu, S. Rosset, R. Tibshirani, and T. J. Hastie. 1-norm support vector machines. In *Advances in neural information processing systems (NIPS)*, pages 49–56, 2004.
- [51] Z. Zhu and A. K. Nandi. Blind modulation classification for mimo systems using expectation maximization. In *Proc. IEEE MILCOM*, pages 754–759, 2014.



Wei Xiong graduated with his PhD degree in the Department of Computer Science, University at Albany. His research focus is on modulation recognition. His work spans a wide range of topics including novel machine learning methods in modulation recognition applications, small-scale cellular networks systems, and network connectivity in resource-limited environments.



Lin Zhang graduated with his Ph.D. in Computer Science at University at Albany—SUNY. His research interests broadly span the areas of data mining, machine learning, and related optimization problems. His current research focuses on two main topics: network analysis and time series analysis.



Maxwell McNeil is a Ph.D. student majoring in computer science at University at Albany. He received his Bachelors of Science in Computer Science and Applied Math (concentrating in Data Analytics) from the University at Albany with a double major in Economics. His research interests include data mining focused on temporal graph mining.



Dr. Petko Bogdanov is an Assistant Professor at the computer science department of University at Albany – SUNY. His research interests include data mining and management and applications to bioinformatics, neuroscience, data-driven nanomaterial design and wireless networks. Previously, he was a postdoctoral fellow at the department of computer science at the University of California, Santa Barbara. He received his PhD and MS in Computer Science from the University of California, Santa Barbara

in 2012 and his BE in Computer Engineering from Technical University—Sofia in 2005. Dr. Bogdanov is a member of the IEEE and the ACM and his research has been supported by grants from NSF, DARPA and ONR. He currently serves as an Associate Editor for IEEE Transactions on Knowledge and Data Engineering (TKDE).



Dr. Mariya Zheleva is an Assistant Professor in the Department of Computer Science at University at Albany – SUNY. Her research is at the intersection of wireless networks and Information and Communication Technology for Development. She has done work on small local cellular networks, data-driven Dynamic Spectrum Access, spectrum management and sensing and network performance and characterization. She is the founder and director of the UbiNet Lab at University at Albany. She holds a PhD and MS

in Computer Science from the University of California, Santa Barbara and a M.Eng. and B.Eng. in Telecommunications from the Technical University – Sofia. Her research is funded by the National Science Foundation and Microsoft.