



Wireless Channel Charting: Theory, Practice, and Applications

Paul Ferrand, Maxime Guillaud, Christoph Studer, Olav Tirkkonen

► To cite this version:

Paul Ferrand, Maxime Guillaud, Christoph Studer, Olav Tirkkonen. Wireless Channel Charting: Theory, Practice, and Applications. IEEE Communications Magazine, 2023, 10.1109/MCOM.001.2200344 . hal-04067052

HAL Id: hal-04067052

<https://inria.hal.science/hal-04067052>

Submitted on 6 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Wireless Channel Charting: Theory, Practice, and Applications

Paul Ferrand, *Senior Member, IEEE*, Maxime Guillaud, *Senior Member, IEEE*,
Christoph Studer, *Senior Member, IEEE*, Olav Tirkkonen, *Fellow, IEEE*

Abstract—Channel charting is a recently proposed framework that applies dimensionality reduction to channel state information (CSI) in wireless systems with the goal of associating a pseudo-position to each mobile user in a low-dimensional space: the *channel chart*. Channel charting summarizes the entire CSI dataset in a self-supervised manner, which opens up a range of applications that are tied to user location. In this article, we introduce the theoretical underpinnings of channel charting and present an overview of recent algorithmic developments and experimental results obtained in the field. We furthermore discuss concrete application examples of channel charting to network- and user-related applications, and we provide a perspective on future developments and challenges as well as the role of channel charting in next-generation wireless networks.

I. INTRODUCTION

MODERN wireless communications systems rely on ever-increasing bandwidth and number of antennas to enable better service. This trend results in a growing number of parameters required to represent the state of the wireless channel, even though channel modeling studies have demonstrated that the wireless channel can be characterized by a comparatively small set of parameters, such as the locations of antennas and environmental scatterers.

Communication protocols and applications increasingly rely on knowledge about the geographical location of the user equipments (UEs) to support high-level network functions, such as handover between base stations and proximity detection [1]. Location-based services are expected to bring a significant part of the value provided by 6G networks [2, Chap. 3]. Unfortunately, existing positioning approaches are either imprecise or require a significant amount of dedicated resources to obtain feedback from the UE [3]. They are also not universally applicable: UEs may not have global navigation satellite system (GNSS) positioning capabilities, they may be indoors or without satellite coverage, or they may forbid sharing the position with a third party.

For many location-based services, absolute position information is in fact not required: a *pseudo-location* that accurately characterizes short-distance neighborhood relationships may be sufficient. Capitalizing on this, channel charting (CC) has recently emerged as an alternative to classical localization techniques that alleviates some of the aforementioned drawbacks [4]. CC consists of learning a mapping between CSI samples and points in a so-called *channel chart*, which is defined as a compact set in a low-dimensional Euclidean space. This mapping is constructed such that nearness between two points on the channel chart indicates similarity between the corresponding CSI samples, their propagation conditions, and

thus their proximity in the physical space. Desirable properties of the channel chart are consistency in time and across the UEs, and low dimensionality.

Every CSI estimate gives rise to a sample represented by a point in the space of CSI values. In order to learn a meaningful mapping between a set of samples in the CSI space and their representations in the channel chart, CC leverages self-supervised dimensionality reduction techniques [5]. Self-supervision is a distinctive characteristic compared to classical positioning techniques, as it allows channel chart generation purely from collected CSI samples—no reference location information from GNSS or from ad-hoc surveys is required. As a consequence of self-supervision, a channel chart pseudo-position only constitutes relative position information w.r.t. other CSI samples and not an absolute position in physical space. This important aspect is illustrated in Fig. 1, which provides a side-by-side comparison between CSI fingerprinting and channel charting. The lack of global geometric consistency in channel charts can be alleviated by resorting to semi-supervised CC approaches, in which a subset of transmit locations are known. Being entirely data-driven, CC naturally handles complex propagation geometries while avoiding the complexity associated with possibly mismatched geometric models.

Many radio resource management (RRM) tasks can be designed to utilize a relative position within the channel chart [6]. The pseudo-position enables infrastructure base stations (BS) to perform context-based RRM, such as SNR prediction or pilot sequence distribution [7], [8]. CC can answer questions such as “are two users close?” or “is the user moving towards an area with poor service?” with minimal protocol overhead, while preserving location privacy and processing only CSI data that is acquired at the BS anyway.

In this article we provide an intuitive idea of the principles of CC (Secs. II and III), and of the related algorithmic and practical implementation aspects (Sec. IV). Since the introduction of CC in [4], several research groups have demonstrated that meaningful channel charts can be practically obtained using data originating from both simulated [4] and measured propagation environments [9]; these will be discussed in Sec. V. We showcase two concrete examples that leverage CC and provide a list of potential applications (Sec. VI), and we discuss challenges, extensions, and future research perspectives (Sec. VII). A regularly updated collection of CC-related publications, code, and datasets may be found on the *Channel Charting Resources* website [10].

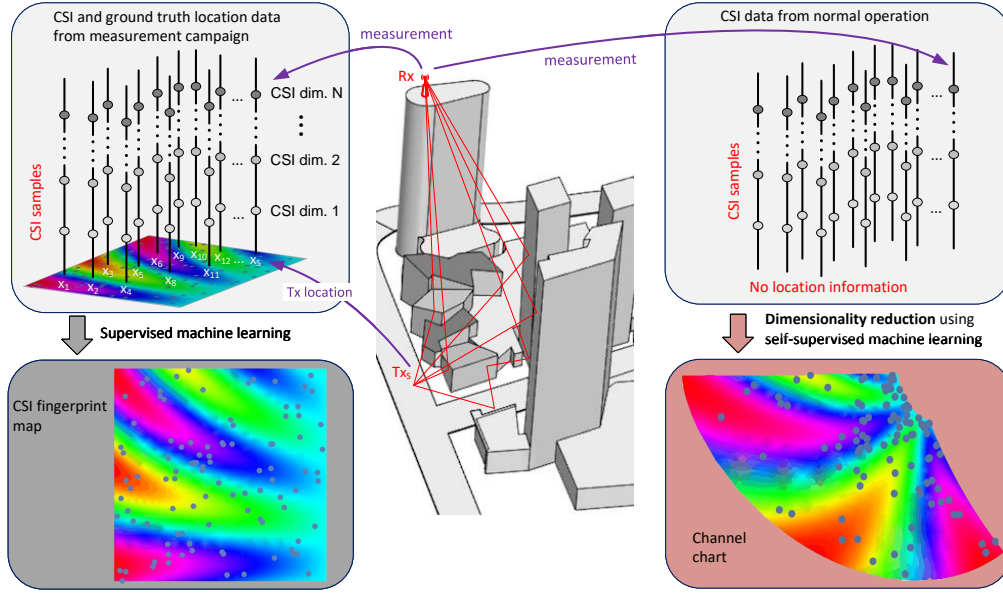


Fig. 1: CSI fingerprinting vs. channel charting (CC). CSI fingerprinting (left) requires CSI samples labeled with location information for training. CC (right) uses unlabeled CSI samples only, thus avoiding costly measurement campaigns; the resulting chart is locally consistent but not necessarily globally accurate.

II. THE WIRELESS CHANNEL MANIFOLD AND ITS SAMPLING

Wireless signals propagate from a transmitting antenna to a receiving antenna according to Maxwell's equations. The set of parameters governing these equations include the positions of the transmitter (Tx) and receiver (Rx) antenna elements, the carrier frequency, as well as the nature and positions of scattering objects in the environment. The effective channel response experienced by a user is a continuous function of these parameters.

Consider, e.g., a set-up with a static BS receiver equipped with an antenna array in a static scattering environment, and a single-antenna mobile user transmitter located anywhere in a two-dimensional (2D) coverage area. The complex-valued vector of channel responses measured at the Rx antennas at a given instant in time (a *CSI sample*) is a point in the *CSI space*, continuously varying as a function of the Tx position. This is the *antenna array manifold* representation of direction-finding algorithms [11]. In our example with static scatterers, the locus of possible CSI samples thus forms a continuous image of the 2D coverage area in the CSI space. Not all CSI samples at the receiver are possible: only those CSI vectors that lie on the image of the coverage area can occur in reality.

This image can be understood as a crumpled, folded, and stretched surface embedded in the high-dimensional CSI space, representing the coverage area. It is not necessarily a manifold—the image may intersect with itself if different locations yield the same CSI. However, such nonuniqueness is not common, and we shall see that it helps to think of the image as a manifold. A CSI sample is a point from this surface. Channel charting is based on the idea that with a sufficiently large number of CSI samples, one may identify the dominant features of the image of the coverage area, and attempt to unfold it into a compact set in a low-dimensional space.

Consider now the situation from a channel modeling perspective. Tractable models use simplified plane-wave approximations of wave equations, where the channel depends on a relatively small subset of the most essential environmental parameters. Absorption, reflection, and diffraction, as well as scattering effects within the environment are typically considered for each wave. While such models postulate a propagation model and attempt to estimate its parameters, CC *jointly* learns the model and its parameters. This key difference enables CC to operate under complex and realistic propagation conditions, affected by multi-path, reflection, diffraction, shadowing, etc. These effects are contained in the data and used in training the channel chart, rather than considered to be corrupting an assumed propagation model.

Reality is of course more complex than a static scattering scenario. However, most objects significantly affecting propagation are static (such as buildings) or have limited mobility (such as cars). In an environment with some mobility, the image of the 2D coverage area would not be a surface, but an object of higher dimension. The relative importance of added dimensions will depend on the impact of the underlying movement on the array manifold.

CSI dimensionality has tremendously increased in recent years. Transmission bandwidths are larger, and base stations and UEs have more antenna elements; contemporary BSs commonly operate with 64 antennas over a 20 MHz bandwidth with around 1200 subcarriers. The dimension of each CSI sample containing all subcarriers and antennas would then easily be in tens of thousands. During normal operation of a radio access network (RAN), the channel from a user may be measured hundreds of times per second. CC summarizes it into a lower-dimensional representation of the data akin to the position of a user on a map. This information can be used as a pseudo-position at higher layers for multiple purposes, as discussed in Sec. VI.




Fig. 2: Example of dimensionality reduction. Top: Data samples around an a priori unknown manifold. Middle: Illustrative "unfolded" manifold. Bottom: An embedding of the dataset obtained by associating 1D coordinates to each sample.

III. MANIFOLD LEARNING VIA DIMENSIONALITY REDUCTION

Dimensionality reduction (DR) techniques are used in CC to (i) identify the CSI manifold within a CSI dataset and (ii) embed it into a low-dimensional chart. Fig. 2 illustrates a simple example in which the data lies on a spiral that is locally similar to a line, i.e., a 1D manifold embedded in 2D space. DR strives to unfold the embedded manifold in the *ambient* (here, 2D) space into its natural representation in the *latent* (here, 1D) space. DR has a rich history in statistical learning [5], where it is used to visualize datasets and to extract meaningful features from high-dimensional data.

DR algorithms may be classified, e.g., by

- whether they produce a parametric or non-parametric mapping of samples in the ambient space to the latent space;
- whether they estimate the intrinsic manifold dimension;
- whether they operate online on a stream of data, or offline on batches of data;
- whether they distinguish between connected components in the underlying manifold;
- whether they output layered embeddings, gradually adding *layers* or *dimensions* until a desired metric is achieved, or not.

One important issue which naturally arises in the dimensionality reduction process when applied to real-world datasets is the presence of noise (due to imperfect CSI estimation) in the input data samples, which can corrupt the manifold learning process. As seen in Fig. 2, due to noise, the latent space is not truly 1D after unfolding. DR algorithms therefore project the data to limit the latent space dimension. Noise may also introduce false neighbors, as exemplified in the figure, which may detrimentally effect the success of the procedure if not handled properly; many DR algorithms rely on creating local

neighborhoods in the ambient space to learn the manifold. Wrongly considering these two samples as neighbors would short circuit two folds of the spiral, and distort the inferred topology of the manifold. Noise also impacts the estimation of the intrinsic dimensionality of the manifold, as can be intuitively understood from Fig. 2. The noise variance may be larger than the variance of the manifold in some area, such that the estimated dimensionality will reflect noise dimensionality.

As DR methods are self-supervised, it is difficult to define metrics for measuring unfolding accuracy. Most metrics rely on comparisons between the unfolded manifold and a reference space, with a known pairing between samples in the manifold and the reference spaces. The reference space can be the ambient space, or the geographical location space if known by supervision. Preservation of key topological or geometrical measures between samples in the latent space and the reference space is evaluated. Common topological metrics are *trustworthiness* and *continuity* [5], which compare the neighborhoods of points between the two spaces. Samples sharing the same neighbors in both spaces will score positively in both metrics; samples being neighbors in one space but not another will affect the metric negatively. A meaningful geometrical metric is the Kruskal stress [12], measuring how well distances between pairs of samples in the latent space match distances in the reference space.

IV. WIRELESS CHANNEL CHARTING

A generic scenario leveraging CC consists of UEs moving in space and a BS implementing DR. The process is typically carried out in two phases.

In the first phase, the BS learns the channel chart. While the mobile UEs are transmitting pilot signals from different locations and time instants, the BS uses the received pilot signals to estimate CSI samples. These CSI estimates are converted into CSI features through a feature extraction function; the resulting CSI features are stored in a database. After constructing the CSI feature dataset, the BS performs dimensionality reduction. For parametric DR, one learns a channel charting function. For nonparametric DR, one directly computes the coordinates of points in the channel chart for every CSI sample. This process is self-supervised, i.e., it does not require any ground-truth location information from the mobile users.

In the second phase, the BS uses the learned channel-charting function or the channel chart positions of the previous samples together with an out-of-sample extension method [13] to map a new CSI sample to the coordinates of a point in the channel chart. This new point is then made available in real time for any chart-based application.

The two-phase approach outlined above does not fully leverage the advantages of self-supervision. The continuous acquisition of CSI in a practical system would enable learning and updating the channel chart continuously. This would ensure that the channel chart is always up-to-date and accurately reflects long-term changes in the environment. The development of such life-long machine learning methods is an active research topic (see Sec. VII).

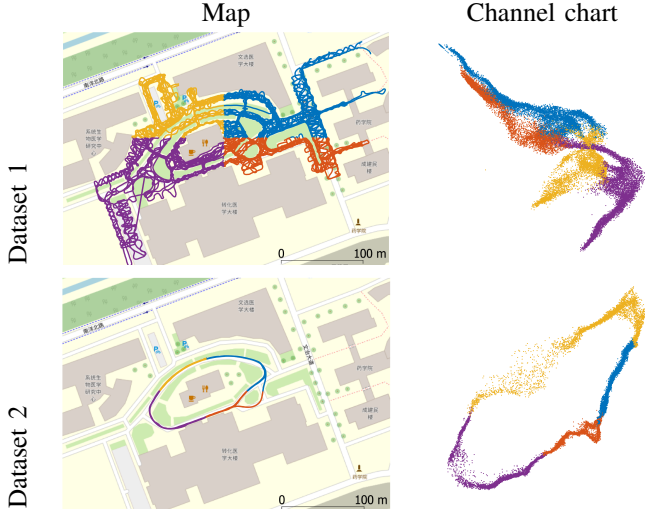


Fig. 3: Channel charts from [9] for a mobile user moving in a campus area. Despite their self-supervised nature, the obtained channel charts (right) are topologically similar to the geographical positions (tracks with matching colors, left).

One of the most critical steps is the extraction of CSI features. The ultimate goal of this step is to prepare the CSI samples such that dimensionality reduction is most effective, meaning that an accurate channel chart can be learned with practical dataset sizes and acceptable complexity. First, CSI feature extraction ensures that only large-scale properties of the propagation environment are captured while small-scale fading and noise are ignored. Effective CSI features may include angle-of-arrival, relative time-of-flight, or frequency selectivity/power-delay profiles. The features should be sufficiently expressive to prevent aliasing caused by nonunique CSI features while unfolding the manifold. Second, CSI feature extraction must mitigate the impact of hardware impairments at both Tx and Rx, taking into account system and hardware components in order to ensure that the features are resilient to timing and synchronization errors, residual carrier frequency and sampling rate offset, phase noise, etc. Methods bypassing feature extraction and directly processing CSI samples have been introduced recently; however, they require significantly larger dataset sizes to deal with variations induced by hardware impairments. Third, CSI feature extraction serves the purpose to reduce the CSI feature dimension with the goals of reducing dataset sizes and accelerating manifold learning.

The other critical step in the first phase is dimensionality reduction. Many methods can be used for this for CC; parametric methods that rely on neural networks (NNs) have shown to be particularly effective [9]. The reasons for their efficiency are as follows: First, NNs are powerful function approximators, which means that an explicit parametric mapping from CSI feature to a point in the channel chart can be learned, even for complicated channel manifolds. They naturally allow for so-called *out-of-sample* extension, whereby a new data point can be mapped to the low-dimensional latent space by means of the learned network. Second, NNs can be efficiently and accurately trained even for large datasets by leveraging stochastic gradient descent. Powerful deep learning frameworks enable easy formulation of

complex network architectures, loss functions, and regularizers, without having to implement DR from scratch. Third, one can readily utilize side information obtained during the acquisition process of CSI sampling to assist dimensionality reduction and improve the learned channel chart. For example, in [9] sample timestamp information is used to improve channel charting. Generally, any proximity information can be used to regularize NN training and significantly improve the accuracy of the learned channel chart [4].

In order for channel charting to be maximally effective, CSI feature extraction and dimensionality reduction should be *jointly* designed, and adapted to the communication system and Tx/Rx hardware.

V. EXPERIMENTAL RESULTS

Channel charting has been successfully applied to both synthetic channels [4], [7], [13]–[15] and real-world CSI measurements [9], using DR algorithms, such as Sammon’s mapping or Laplacian eigenmaps (see [5]), as well as NN approaches, such as Siamese or triplet networks.

Fig. 3 depicts two channel charts from [9] obtained by training a triplet-based NN on real-world CSI measurements. The datasets were acquired using a commercial 4G base station, with a rooftop antenna array with 32 dual-polarized antenna elements in a uniform 8×4 rectangular array. The UE has a single antenna. Orthogonal frequency division multiplexing with 288 subcarriers in a 10 MHz bandwidth at 2.5 GHz is used. Dataset 1 includes some 3 hours of pedestrian movement—at about 1 m/s—around the dense urban coverage area, with both line-of-sight (LoS) and non-LoS areas, amounting to several kilometers worth of tracks. The resulting channel chart shows excellent connectivity of the four colored subsets, but also illustrates the inherent uncertainty about rotations and mirroring within a channel chart when compared to the geographical position. Dataset 2 consists in a 20-minute run where the UE was moved multiple times along an approximately elliptic trajectory around the square. The resulting channel chart correctly “connects” the loop as the meaningful representation of the CSI dataset. These results show that channel charting is able to recover topological properties from measured data in a completely self-supervised manner. Both channel charts in Fig. 3 score well with respect to the metrics discussed in Sec. III. A complete numerical performance assessment is available in [9].

VI. APPLICATIONS

A channel chart, being merely a low-dimensional representation of the state of the wireless channel, is not a goal in itself. Instead, it should be seen as a versatile platform upon which applications can rely to infer, predict, and process propagation-related information. We now showcase two example applications that demonstrate the efficacy of CC.

The first example deals with mmWave beam discovery in a dual-band set-up. A channel chart is constructed for Dataset 1 from Fig. 3, for the massive MIMO CSI. The associated cell connectivity for the mmWave band is simulated as depicted in Fig. 4. CSI samples in the channel chart corresponding to the training set are colored according to their proximity with

TABLE I: Applications of Channel Charting

Application	Details	Leveraged channel chart property
Radio resource management	Handover prediction [10], pilot allocation [8], rate or protocol adaptation, predictive buffering [1], BS association.	Temporal consistency, user trajectory extrapolation
mmWave beam management	Accelerated beam discovery [7], predictive beam tracking	Temporal and spatial consistency, multi-view chart
Device grouping for D2D	Identify and cluster nearby mobile devices [13]	Spatial consistency
CSI compression	Representing CSI using the minimal number of dimensions, for efficient quantization and feedback	Data-driven, time consistency
Proximity detection, geofencing	Detect user proximity or relative position to a fixed feature on the ground (shop, prohibited areas) using the channel chart	Spatial consistency
Network event labeling	Identify and cluster link failures events and low performance spots for network planning purposes	Spatial and temporal consistency
Context awareness	Infer contextual information (e.g., pedestrian or in-vehicle; work commute or unusual trip) from the radio link	Data driven, spatial and temporal consistency
Localization	Supervised or semi-supervised version of charting [10]	Spatial and temporal consistency
Channel modeling and simulation	Understand and model the wireless propagation medium, generate realistic CSI samples	Data-driven, temporal consistency

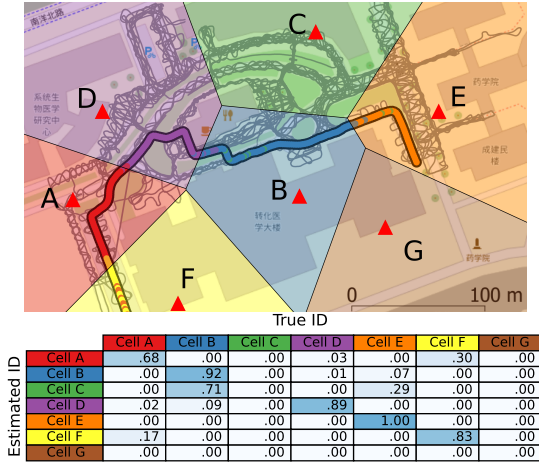


Fig. 4: Example of chart-based cell association in a network of seven mmWave cell. (Top) test user trajectory colored according to the estimated cell ID. The channel chart (not shown) was constructed from Dataset 1, shown in gray. (Bottom) confusion matrix showing the probability of the true cell ID, for each estimated cell ID.

mmWave BSs (red triangles). Samples from the test dataset (trajectory depicted in the figure) are mapped to the channel chart, and their mmWave cell association is estimated from the color of the neighboring training set points in the channel chart. This approach achieves 85% accuracy; most errors occur at cell-edge locations where a neighboring cell is chosen instead.




Fig. 5: ROC curves for channel-chart-based proximity detection.

The second example deals with proximity detection. We train two channel charts using Dataset 1 again: a pseudo-location based fully self-supervised channel chart and a semi-supervised version using known Tx-position information acquired every five minutes. Test user positions are sampled from the channel chart, and their position is averaged over five seconds. Pairs of test users are declared neighbors if their distance in the channel chart is lower than some threshold and non-neighbors otherwise. The receiver operating characteristic (ROC) curves of this hypothesis test curves depicted in Fig. 5 indicate that it is possible to achieve a high true positive detection rate with few false positives. For example, a 10% false positive rate yields a true positive rate of 96% for the self-supervised and 98% for the semi-supervised case. Pseudo-positioning via CC can thus replace positioning with minimal accuracy loss in this proximity detection example.

Beyond the two applications above, Table I lists a range of potential applications of CC and identifies the relevant channel chart properties. Several of the applications in this table can be solved without resorting to a channel chart. For instance, the correspondence between mmWave beam direction and the CSI measured on a low-frequency band can be learned directly. However, a channel chart provides a simple unified representation of propagation-related characteristics which may be shared among several applications using a standardized interface, thus allowing significant savings in computation and storage over parallel end-to-end solutions each involving a separate learning problem. Furthermore, the CC-based solution does not require directly exposing CSI or proprietary base-station hardware details. In this “*channel charting as a service*” paradigm, the computationally-intensive channel chart update and maintenance tasks are delegated to the host network, while simple but powerful access to propagation-related characteristics can be extended up to the application layer, enabling the ex-post development of innovative applications not listed in Table I.

VII. CHALLENGES, EXTENSIONS, AND FUTURE DIRECTIONS

A successful implementation of channel charting presents unique challenges beyond solving the DR problem.

1) *Continuous learning*: The channel chart should evolve over time, reflecting changes in the propagation environment. Some changes may be permanent (e.g., construction of a new building) and will require a channel chart update, while others might be temporary, for instance the presence a large vehicle altering propagation. In the former case, past information shall be forgotten, while in the latter case, the presence/absence of the vehicle might be captured as extra dimension of the channel chart. It would then constitute information that could be inferred from CSI and used for context-aware applications. Alternatively, the channel chart can be made robust to such changes, considering them a nuisance rather than useful information.

2) *Online and real-time aspects*: CSI typically streams in continuously at a rate of hundreds samples per second. This situation departs from classical machine learning problems where fixed-size datasets are processed offline. Algorithmic adaptations are required to achieve online, life-long learning. In CC applications, both learning (charting) and inference needs to be performed with stringent real-time and storage constraints, which opens up numerous challenges related to the real-time implementation of machine learning algorithms for the physical layer.

3) *Distributed implementation*: In modern radio access networks, mobile devices can be served by multiple physically distinct transmission points (TRPs) and in multiple frequency bands. In this context, the capability of constructing joint or distributed channel charts across multiple TRPs and bands, i.e., *multi-view* channel charts [15] will be a key enabler for handover and beam steering applications.

4) *Array geometry and propagation scenario*: CC for indoor applications presents unique challenges: the keyhole effect typically observed in outdoor-to-indoor propagation scenarios might limit the expressivity of the channel chart, while the relevance of charting for indoors systems will hinge on a favorable geometry of the antenna system (e.g., data gathered from a distributed antenna system will capture a richer range of environmental properties than a single array with co-located antennas). Analog Tx- and/or Rx-side beamformers are commonly used during CSI estimation, and they are required in high frequency bands in order to guarantee a sufficient link budget—this yields a partial channel observation that complicates channel charting. A proper framework for handling time-varying beamformers within the charting toolchain is necessary [7]. In this context, closed-loop strategies (such as reinforcement learning approaches) involving dynamic adjustments of the CSI measurement process can potentially further enhance the chart coverage and accuracy.

5) *Privacy aspects*: Some of the applications described in Table I (e.g., proximity detection) reveal that the channel chart's pseudo-position information may have to be considered under the angle of privacy protection. Note that this is generally true for RANs which have provided some kind of localization services since 2G. With the increasing importance of proper handling of private information, it is expected that 6G will have built-in technical features guaranteeing privacy through a powerful data governance architecture for better societal acceptance [2, Chap. 30].

REFERENCES

- [1] M. Kasparick, R. L. G. Cavalcante, S. Valentin, S. Stańczak, and M. Yukawa, "Kernel-based adaptive online reconstruction of coverage maps with side information," *IEEE Trans. Veh.*, vol. 65, no. 7, pp. 5461–5473, Jul. 2016.
- [2] W. Tong and P. Zhu, Eds., *6G: The Next Horizon: From Connected People and Things to Connected Intelligence*. Cambridge University Press, 2021.
- [3] O. Alamu, B. Iyaomolere, and A. Abdulrahman, "An overview of massive MIMO localization techniques in wireless cellular networks: Recent advances and outlook," *Elsevier Ad Hoc Netw.*, vol. 111, Feb. 2021.
- [4] C. Studer, S. Medjkouh, E. Gönültaş, T. Goldstein, and O. Tirkkonen, "Channel charting: Locating users within the radio environment using channel state information," *IEEE Access*, vol. 6, pp. 47 682–47 698, Aug. 2018.
- [5] L. Van Der Maaten, E. Postma, and J. Van den Herik, "Dimensionality reduction: a comparative review," *J. Mach. Learn. Res.*, vol. 13, no. 10, pp. 66–71, Oct. 2009.
- [6] A. Bourdoux & al., "6G white paper on localization and sensing," University of Oulu, Tech. Rep., Jun. 2020, ISBN 978-952-62-2674-3.
- [7] P. Kazemi, T. Ponnada, H. Al-Tous, Y.-C. Liang, and O. Tirkkonen, "Channel charting based beam SNR prediction," in *Proc. Eur. Conf. on Netw. Com.*, Porto, Portugal, Jun. 2021.
- [8] L. Ribeiro, M. Leinonen, D. Djelouat, and M. Juntti, "Channel charting for pilot reuse in mMTC with spatially correlated MIMO channels," in *Intl. Conf. Cognitive Radio Oriented Wireless Netw.*, Taipei, Taiwan, Dec. 2020.
- [9] P. Ferrand, A. Decurninge, L. G. Ordoñez, and M. Guillaud, "Triplet-based wireless channel charting: Architecture and experiments," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 8, pp. 2361–2373, Aug. 2021.
- [10] "Channel charting resources," website <https://channelcharting.github.io/>.
- [11] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Ant. and Propag.*, vol. 34, no. 3, pp. 276–280, Mar. 1986.
- [12] J. B. Kruskal, "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis," *Psychometrika*, vol. 29, no. 1, Mar. 1964.
- [13] H. Al-Tous, T. Ponnada, C. Studer, and O. Tirkkonen, "Multipoint channel charting-based radio resource management for V2V communications," *EURASIP J. Wireless Commun. Netw.*, Dec. 2020.
- [14] P. Agostini, Z. Utkovski, and S. Stańczak, "Channel charting: An Euclidean distance matrix completion perspective," in *IEEE Intl. Conf. Acoust., Speech, Signal Process.*, Barcelona, Spain, May 2020.
- [15] J. Deng, S. Medjkouh, N. Malm, O. Tirkkonen, and C. Studer, "Multipoint channel charting for wireless networks," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, USA, Oct. 2018.

PLACE
PHOTO
HERE

Paul Ferrand received the M.Sc. degree in computer science from the Université de Lyon, France, in 2009, and the Ph.D. degree from INSA Lyon, France, in 2013. He was a Post-Doctoral Fellow with INRIA, France, where he was involved in energy-efficient communications. From 2014 to 2021, he was a Senior Researcher at the Huawei Research Center, Paris, France. His main research interests lie within the field of signal and information processing for the physical layer of wireless communications.

PLACE
PHOTO
HERE

Maxime Guillaud is a senior researcher at Inria in Lyon. He has over 20 years of expertise in the domain of wireless communications in both academic and industrial research environments. He received his Ph.D. degree in electrical engineering and communications from EURECOM and Telecom ParisTech in 2005. From 2006 to 2010, he was a Senior Researcher with FTW. From 2010 to 2014, he was a research associate at Vienna University of Technology. From 2014 to 2022, he was with Huawei Technologies. His expertise lies the physical layer of

RANs, including transceiver algorithms, channel modeling, and modulation design.

PLACE
PHOTO
HERE

Christoph Studer is an Associate Professor at the Department of Information Technology and Electrical Engineering at ETH Zurich in Switzerland. Dr. Studer received his M.S. and Ph.D. degrees in electrical engineering from ETH Zurich in 2005 and 2009, respectively. Between 2009 and 2013, he held postdoctoral positions at ETH Zurich and Rice University in Houston, TX. In 2014, he joined Cornell University, Ithaca, NY, as an Assistant Professor. From 2019 to 2020, he was an Associate Professor at Cornell University and Cornell Tech in New York

City. In 2020, he joined ETH Zurich. Dr. Studer's research interests include the design of digital integrated circuits, as well as wireless communications, digital signal processing, numerical optimization, and machine learning.

PLACE
PHOTO
HERE

Olav Tirkkonen is a professor in communication theory at Aalto University, Finland, where he has held a faculty position since 2006. He received his M.Sc. and Ph.D. degrees in theoretical physics from Helsinki University of Technology in 1990 and 1994, respectively. He held post-doctoral positions at UBC, Vancouver, Canada, and NORDITA, Copenhagen, Denmark, and was with Nokia Research Center, Helsinki, Finland from 1999 to 2010. In 2016–2017, he was Visiting Associate Professor at Cornell University, Ithaca, NY, USA. His current research interests

are in coding for random access and quantization, quantum computation, and machine learning for cellular networks.