Self-Supervised Learning for Early Preamble Collision Detection in Cellular IoT Networks

Daniela M. Casas-Velasco[1], **Diogo M. da Cunha**[1], **Marco A. Guerra**[1], **Giancarlo M. Cardenas**[1], **Carlos A. Astudillo**[1], **Nelson L. S. da Fonseca**[1]

[1] Institute of Computing – University of Campinas (Unicamp)
Campinas – SP – Brazil

{d211826, d249418, m184714,g262946}@dac.unicamp.br

castudillo@unicamp.br,nfonseca@ic.unicamp.br

Abstract. *Massive Machine-Type Communications (mMTC) scenarios in Cellular Internet of Things (CIoT) networks may generate peaks of simultaneous access attempts, making the random access (RA) procedure a critical bottleneck. In such conditions, preamble collisions increase contention, retransmissions, and access delay. Recent supervised approaches based on power-delay-profile (PDP) features have shown promising performance for early preamble collision detection, but they depend on labeled datasets. In this domain, labeled data is costly and may not be available (e.g., in new network deployments) because collision labels require detailed simulation and protocol-level verification. This paper investigates self-supervised learning (SSL) to reduce label dependence while improving downstream collision classification and cross-scenario robustness. A Transformer-based encoder is pre-trained on unlabeled PDP sequences, and representations learned are then transferred via fine-tuning with a small labeled subset, employing a simple multilayer perceptron (MLP) classifier as a prediction head. Two masking-based SSL pretext tasks are evaluated, namely random token masking and contiguous masked-span prediction, using the datasets with varying characteristics, such as channel type, speed, and coverage radius. Simulation results show that SSL improves the downstream task performance in all same-scenario evaluations and in five of the six cross-scenario transfer pairs, with the largest gains under domain mismatch. These findings indicate that SSL reduces dependence on labeled data while improving transferability across different radio channel conditions.*

Resumo. *Cenários de Massive Machine-Type Communications (mMTC) em redes Cellular Internet of Things (CIoT) podem gerar picos de tentativas simultâneas de acesso, tornando o procedimento de acesso aleatório (random access, RA) um gargalo crítico. Nessas condições, colisões de preâmbulos aumentam a contenção, as retransmissões e o atraso de acesso. Abordagens supervisionadas recentes baseadas em características do power-delay profile (PDP) têm apresentado resultados promissores para detecção antecipada de colisão de preâmbulos, mas dependem de conjuntos de dados rotulados. Nesse domínio, dados rotulados são custosos e podem não estar disponíveis (e.g., em novas implantações de rede), uma vez que os rótulos de colisão exigem simulação de-*

talhada e verificação em nível de protocolo. Este artigo investiga o uso de self-supervised learning (SSL) para reduzir a dependência de rótulos e, ao mesmo tempo, melhorar a classificação de colisão na tarefa downstream e a robustez entre cenários. Um codificador baseado em Transformer é pré-treinado com sequências de PDP não rotuladas, e as representações aprendidas são então transferidas por meio de fine-tuning com um pequeno subconjunto rotulado, empregando um classificador simples do tipo multilayer perceptron (MLP) como cabeça de predição. Duas tarefas de pré-texto baseadas em mascaramento são avaliadas, a saber, random token masking e contiguous masked-span prediction, utilizando conjuntos de dados com características distintas, como tipo de canal, velocidade e raio de cobertura. Os resultados de simulação mostram que o SSL melhora o desempenho da tarefa downstream em todas as avaliações no mesmo cenário e em cinco dos seis pares de transferência entre cenários, com os maiores ganhos observados sob desajuste de domínio. Esses resultados indicam que o SSL reduz a dependência de dados rotulados ao mesmo tempo em que melhora a transferibilidade entre diferentes condições de canal de rádio.

1. Introduction

Massive Machine-Type Communications (mMTC) is one of the key service categories in modern cellular networks, aiming to support the connectivity of a very large number of devices with heterogeneous traffic profiles. This connectivity is provided by Cellular Internet of Things (CIoT) technologies such as NR, LTE-M, and NB-IoT. In cellular networks with massive access, large peaks of simultaneous random access attempts may be generated. Such a situation occurs after alarm-triggered reports, synchronized monitoring cycles, or high-density events in crowded environments, where a large number of terminals contend for access within a short time interval. Under these conditions, the random access (RA) procedure becomes a critical bottleneck [3GPP 2011, Wei et al. 2015, de Andrade et al. 2017].

In the RA procedure, user equipment (UE) devices initiate access by transmitting a preamble in the Physical Random Access Channel (PRACH). When multiple UEs select the same preamble in the same random access opportunity, collisions occur, increasing contention, retransmissions, and access delay [3GPP 2011, Wei et al. 2015, de Andrade et al. 2017]. These effects are especially harmful in dense CIoT deployments, where fast and reliable access establishment is required despite limited radio resources. As a result, early identification of preamble collisions is an important step toward improving access efficiency and reducing the impact of overload conditions [Cardenas et al. 2025]. By identifying collisions early at the physical layer, the network can avoid the unnecessary allocation of resources for Msg2 (Random Access Response) and subsequent steps, thereby reducing signaling overhead and energy consumption in power-constrained IoT devices. This improvement in access efficiency is crucial for maintaining connectivity and quality of service in massive deployment scenarios.

Recent studies have shown that machine learning can support early collision detection by exploiting physical-layer information available immediately after Msg1 reception (*i.e.*, after preamble detection). In particular, supervised approaches based on power-delay-profile (PDP) features have demonstrated that neural-network-based models can learn discriminative collision patterns and can generalize across different propagation conditions better than several classical classifiers [Jang et al. 2021, Magrin et al. 2019, Yin et al. 2023, Cardenas et al. 2025]. A recent baseline study generated realistic datasets spanning different channel, mobility, and coverage conditions, and reported strong performance for supervised neural models under both same-scenario and cross-scenario evaluation [Cardenas et al. 2025].

Despite these promising results, supervised collision detection still depends on labeled datasets whose construction is costly and time-consuming. In the baseline setting, labels are not directly observable from the received PRACH signal and must be obtained through detailed simulation of the RA procedure, including Msg3-based collision verification [Cardenas et al. 2025]. In addition, representative collision events are difficult to obtain through conventional data collection, since informative collision patterns are scarce and strongly dependent on traffic conditions, user synchronization, coverage radius, and channel characteristics. Consequently, assembling sufficiently diverse labeled datasets for new scenarios becomes burdensome, which limits the scalability of purely supervised approaches.

Self-supervised learning (SSL) offers an attractive alternative because it enables robust representation learning from unlabeled data by defining supervision directly from the input itself. The masking strategy has become one of the most widely adopted SSL paradigms, first applied in natural language processing and later in time-series and structured-data domains [Devlin et al. 2019, Li et al. 2023b, Dong et al. 2023, Hajiramezanali et al. 2023]. In the present context, SSL has the potential to exploit large volumes of unlabeled PDP observations, avoiding the need for collision labels, which are expensive, difficult to obtain at scale, and sometimes unavailable.

Motivated by this gap, this paper investigates whether SSL pre-training can improve downstream collision classification while reducing dependence on labeled data. To this end, two masking-based SSL pretext tasks are evaluated, namely random token masking and contiguous masked-span prediction, using datasets with varying characteristics, such as channel type, speed, and coverage radius, under multiple combinations of pre-training, fine-tuning, and evaluation scenarios.

The results show that SSL consistently improves the downstream classifier under severe label scarcity. In particular, SSL improves all same-scenario evaluations and five of the six cross-scenario transfer pairs. Moreover, the results indicate that no single masking strategy dominates all conditions: contiguous masked-span prediction is especially effective in several transfer settings, whereas random masking combined with mixed-domain pre-training yields the strongest average cross-scenario performance.

The main contributions of this work are twofold. First, it introduces an SSL-based pipeline for early PRACH preamble collision detection that exploits unlabeled PDP samples during pre-training and transfers the learned encoder to the downstream classification task. Second, it provides a comparative analysis of two masking-based SSL objectives across multiple combinations of pre-training, fine-tuning, and evaluation scenarios. Results show that SSL is particularly beneficial under cross-scenario transfer learning, where label scarcity and domain mismatch are most critical.

The remainder of this paper is organized as follows. Section 2 reviews the related work on preamble collision mitigation and self-supervised learning. Section 3 describes the datasets used in our study. Section 4 presents the proposed self-supervised pipeline, including the masking strategies and the transfer learning process. Section 5 discusses the experimental results for both intra-scenario and cross-scenario evaluations. Finally, Section 6 concludes the paper and suggests future research directions.

2. Related Work

Preamble collision mitigation in CIoT and mMTC has been investigated from different perspectives. Conventional approaches include protocol modifications, tagged preambles, and time-based discrimination mechanisms such as time-of-arrival (ToA), time advance (TA), and successive interference cancellation (SIC) [Cardenas et al. 2025]. Although these methods can improve collision resolution, they often require additional signaling, modifications to the protocol stack, restrictive assumptions on propagation conditions, or increased computational complexity in dense and dynamic scenarios [Cardenas et al. 2025].

Machine-learning-based solutions have emerged as a more flexible alternative. Jang *et al.* proposed a deep-learning-based cellular random access framework capable of learning collision-related patterns directly from signal observations [Jang et al. 2021]. Magrin *et al.* investigated machine learning for LTE RACH collision multiplicity detection [Magrin et al. 2019], while Yin *et al.* studied learning-based preamble collision detection from physical-layer features [Yin et al. 2023]. More recently, Maldonado Cardenas *et al.* developed a supervised framework for early preamble collision detection based on PDP bins and evaluated it under datasets with distinct propagation, mobility, and spatial characteristics [Cardenas et al. 2025]. These studies confirm the potential of machine learning for RA enhancement, but they also reveal a common limitation: their dependence on labeled training data and the performance degradation observed under cross-scenario evaluation.

In parallel, SSL has become an important paradigm for learning representations from unlabeled data. The success of masked prediction in BERT established the effectiveness of reconstructing hidden input elements from context [Devlin et al. 2019]. Similar principles were later extended to time-series analysis, where masked reconstruction has been shown to improve downstream prediction and classification tasks [Li et al. 2023b,Dong et al. 2023]. SSL has also been adapted to structured and tabular data, where meaningful augmentations are harder to define but unlabeled representation learning remains valuable [Hajiramezanali et al. 2023]. In these contexts, masked reconstruction encourages the model to capture the inherent temporal and structural correlations within the data, which is particularly suitable for PDP observations where the multipath components follow specific physical propagation patterns.

In wireless communications, SSL has already been explored for tasks with limited labeled data. Milosheski *et al.* applied SSL to clustering wireless spectrum activity from unlabeled measurements [Milosheski et al. 2023], and Li *et al.* used self-supervised contrastive learning to reduce label requirements in spectrum sensing [Li et al. 2023a]. More recently, Enoki and Astudillo investigated BYOL-based SSL pre-training for spectrogram-based identification of mobile communication technologies, showing that models pretrained on RF spectrograms outperform training from scratch and remain particularly effective in scarce-label regimes [Enoki and Astudillo 2025]. These studies support the premise that SSL is well suited to wireless domains in which labels are expensive, noisy, or unavailable.

However, the use of SSL techniques (in particular masking-based ones) for early PRACH preamble collision detection using PDP bins under datasets that vary in channel conditions, mobility, and coverage range is missing in the literature. This paper addresses this gap by investigating masking-based SSL pre-training for downstream collision classification and by systematically comparing multiple pre-training, fine-tuning, and evaluation scenario combinations. In particular, it assesses whether SSL improves the performance of a downstream MLP under severe label scarcity and whether the learned representations transfer robustly across scenarios with different channel conditions.

3. Dataset Description

The present work uses the same three datasets adopted in the supervised baseline for early preamble collision detection in the random access procedure, namely EPA500, EPA790, and ETU790 [Cardenas et al. 2025]. Such datasets were originally generated through simulation with the MATLAB LTE System Toolbox under realistic Cellular IoT random access conditions, since public datasets containing detailed PRACH signal observations and collision outcomes are not readily available [Cardenas et al. 2025]. In the original data-generation process, the scenarios differ in propagation model and cell radius: EPA500 corresponds to the Extended Pedestrian A (EPA) channel with 500m cell radius, EPA790 uses the same channel model with 790m radius, and ETU790 adopts the Extended Typical Urban (ETU) channel with 790m radius [Cardenas et al. 2025]. This set of scenarios makes it possible to evaluate both matched-domain and mismatched-domain learning conditions.

The original simulation framework considered a large-scale access environment with multiple UEs competing for random access opportunities. For each event, active UEs selected preambles, propagated through the channel, and generated received PRACH signals at the base station. These signals were processed to obtain the corresponding power delay profile (PDP), which was then used as the basis for collision detection [Cardenas et al. 2025]. Since collision labels are not directly observable from Msg1 alone, the baseline dataset construction extended the simulation to later stages of the random access procedure, including Msg3, in order to determine whether a detected preamble effectively corresponded to a successful access or to a collision event [Cardenas et al. 2025]. This labeling requirement is precisely the main practical limitation addressed in this paper, as it makes the generation of annotated datasets costly and time-consuming.

Each sample used in the baseline study corresponds to one PDP bin extracted from the received PRACH signal. More specifically, the full PDP is partitioned into fixed-size bins, and each instance is represented by the 24 amplitude values associated with one bin, together with a binary target indicating collision or non-collision [Cardenas et al. 2025]. In the preprocessed files used in the present work, only these 24 PDP values and the corresponding binary label are retained. Therefore, the input to all models is a 24-dimensional vector, and no additional metadata from the original simulation is used during training or inference. This representation is consistent with the objective of assessing whether local PDP patterns are sufficient for early collision identification.

For experimental purposes, each scenario is provided with separate training and test partitions, following the same organization adopted in the baseline work. The training partition is used in two different ways in this paper. First, most of its samples are employed without labels during the self-supervised pre-training stage. Second, a small labeled subset is reserved for the downstream fine-tuning stage. The test partition remains untouched during model development and is used only for final evaluation. In addition to the three original scenarios, a mixed pre-training set is also considered by combining training samples from EPA500, EPA790, and ETU790. This mixed set is used to investigate whether a more heterogeneous unlabeled pool improves the robustness of the learned representations under cross-scenario transfer.

The use of EPA500, EPA790, and ETU790 is especially suitable for the objectives of this paper. EPA500 and EPA790 share the same propagation model but differ in cell radius, which allows the effect of spatial-domain variation to be analyzed. ETU790, in turn, introduces a distinct channel condition with higher propagation complexity, thereby representing a more challenging transfer setting. As a result, the adopted datasets support both the evaluation of self-supervised learning under matched conditions and the analysis of its generalization ability when the pre-training, fine-tuning, and evaluation domains are not the same. To ensure the statistical relevance of the results, the training and testing partitions were designed to maintain the original class distribution of the baseline study [Cardenas et al. 2025]. Although the pre-training uses a vast amount of data, the final performance metrics are calculated over a fixed test set to ensure that the gains observed are due to the learned representations and not to data overlap.

4. Methodology

4.1. Overview

The proposed method follows a two-stage learning pipeline. In the first stage, a self-supervised model is pre-trained on unlabeled PDP sequences to learn informative signal representations without using collision labels. In the second stage, the learned encoder is transferred to the downstream collision-classification task, where only a small labeled subset is used for supervised training. For comparison purposes, a purely supervised MLP baseline is also trained directly on labeled PDP samples. This setup makes it possible to evaluate whether self-supervised pre-training improves downstream classification under severe label scarcity.

Each input sample is represented by a 24-dimensional PDP vector,

$$\mathbf{x} = [x_1, x_2, \ldots, x_{24}] \in \mathbb{R}^{24}, \tag{1}$$

and, for the downstream task, an associated binary label

$$y \in \{0, 1\}, \tag{2}$$

where $y = 0$ denotes non-collision and $y = 1$ denotes collision. As discussed in Section 3, the adopted datasets correspond to the same PDP-bin representation introduced in the supervised baseline study [Cardenas et al. 2025].

During self-supervised pre-training, labels are discarded and only the PDP vectors are used. Let $\mathcal{D}_u^{(s_{pt})}$ denote the unlabeled pre-training set associated with scenario

$s_{pt} \in \{\text{EPA500}, \text{EPA790}, \text{ETU790}, \text{MIXED}\}$. The mixed setting is obtained by combining unlabeled training samples from the three original scenarios. For downstream training, a small labeled subset $\mathcal{D}_l^{(s_{ft})}$ is used, where $s_{ft} \in \{\text{EPA500}, \text{EPA790}, \text{ETU790}\}$. The test set $\mathcal{D}_{test}^{(s_{ev})}$ is kept separate and is used only for final evaluation.

The implemented split strongly favors the unlabeled stage. Specifically, the pre-training stage uses 99.99% of the available training partition, while only a very small fraction is reserved for the downstream labeled stage. This configuration is intentional, since the main objective is to evaluate SSL under severe label scarcity. The experimental protocol is systematically evaluated using these three scenarios to form different learning setups. Specifically, s_{pt} defines the domain where the model learns its initial features, s_{ft} is the domain used to adapt the classifier head with minimal labels, and s_{ev} is the target domain where the final performance is measured.

4.2. Supervised Baseline

For comparison, a supervised baseline is trained directly on labeled PDP samples. The baseline classifier is a shallow MLP composed of one hidden layer with 48 neurons, ReLU activation, dropout with rate 0.2, and a two-neuron output layer for binary classification. Let $f_b(\mathbf{x})$ denote the logits produced by the baseline model. The classifier is trained with the cross-entropy loss

$$\mathcal{L}_{cls} = -\sum_{c\in\{0,1\}} y_c \log \hat{y}_c, \tag{3}$$

where $\hat{y}_c$ is the posterior probability assigned to class c.

This specific baseline architecture was chosen to match the complexity of the supervised models previously reported in the literature [Cardenas et al. 2025], ensuring a fair comparison. The shallow structure (one hidden layer) is intended to prevent overfitting given the severe label scarcity (only 0.01% of the training set) used during the supervised phase. This baseline provides the reference point against which the SSL-based models are compared.

4.3. Self-Supervised Encoder

The self-supervised encoder is based on a Transformer architecture [Vaswani et al. 2017]. Since each PDP input is a short one-dimensional sequence of length 24, each scalar value is first projected to a higher-dimensional embedding space. More precisely, each element x_i is mapped from $\mathbb{R}$ to $\mathbb{R}^{256}$ through a learnable linear projection. A learnable positional embedding is then added to preserve the order of the PDP samples. The resulting embedded sequence is processed by a Transformer encoder with one layer, eight attention heads, and feedforward dimension 512.

Let $E_\theta(\cdot)$ denote the Transformer encoder and let $P_\theta(\cdot)$ denote the prediction head. Given a masked input $\tilde{\mathbf{x}}$, the pre-training model outputs a reconstruction

$$\hat{\mathbf{x}} = P_\theta(E_\theta(\tilde{\mathbf{x}})). \tag{4}$$

The prediction head is implemented as a two-layer perceptron with dimensions $256 \rightarrow 128 \rightarrow 1$, applied position-wise to the encoded sequence. Pre-training minimizes the mean squared reconstruction error with respect to the original PDP vector,

$$\mathcal{L}_{ssl} = \frac{1}{24}\sum_{i=1}^{24}(\hat{x}_i - x_i)^2. \tag{5}$$

This formulation allows the encoder to learn contextual dependencies among PDP samples without requiring collision labels.

4.4. Masking Strategies for SSL

We considered two masking-based pretext tasks, namely, Random token masking and Contiguous masked-span prediction.

In the Random token masking, six positions are selected uniformly at random from the 24-sample PDP vector and replaced by a fixed mask value. The model then receives the corrupted sequence $\tilde{\mathbf{x}}$ and is trained to reconstruct the original signal. This strategy follows the general principle of masked prediction used in self-supervised learning [Devlin et al. 2019, Li et al. 2023b, Dong et al. 2023], while adapting it to short PDP sequences.

In Contiguous masked-span prediction, a contiguous span of six consecutive positions is masked. The model then receives the corrupted PDP vector and is trained to infer the missing span from the surrounding context. In contrast to random token masking, this strategy removes a localized region of the input, encouraging the encoder to capture short-range structural dependencies within the PDP sequence.

The comparison between the two pretext tasks makes it possible to assess whether localized masking is more effective than dispersed masking for learning discriminative representations of PDP structure.

4.5. Transfer to the Downstream Task

After self-supervised pre-training, the encoder is transferred to the downstream collision-classification stage. Only the encoder is retained, while the reconstruction head is discarded. The encoder parameters are then frozen, and the sequence representation is obtained by mean pooling over the 24 encoded positions:

$$\mathbf{z} = \frac{1}{24}\sum_{i=1}^{24}\mathbf{h}_i, \tag{6}$$

where $\mathbf{h}_i \in \mathbb{R}^{256}$ is the encoded representation of the i-th PDP position. The pooled feature vector $\mathbf{z}$ is passed to a classification head with architecture $256 \to 64 \to 2$, including ReLU activation and dropout.

Thus, for the SSL-based models, the downstream classifier is defined as

$$\hat{\mathbf{y}} = C_\phi(\mathbf{z}), \tag{7}$$

where $C_\phi(\cdot)$ is the task-specific classification head. The downstream stage is trained with the same cross-entropy loss used for the baseline model.

Freezing the encoder is consistent with the objective of evaluating the quality of the representations learned during self-supervised pre-training. In this way, improvements in downstream performance can be attributed primarily to the learned representation rather than to extensive end-to-end supervised adaptation.

4.6. Training Configuration

All models were implemented in PyTorch and trained on Google Colab T4 GPUs, with the random seed fixed to 42. In the self-supervised stage, the batch size is set to 2048, the learning rate to 10^{-3}, the maximum number of epochs to 20, and the early-stopping patience to 3. In the downstream stage, the batch size is 16, the learning rate is 10^{-3}, the maximum number of epochs is 500, and the patience is set to 10. In both stages, optimization is performed with Adam optimizer.

The complete experimental protocol is defined by the combination of three factors: pre-training scenario s_{pt}, fine-tuning scenario s_{ft}, and evaluation scenario s_{ev}. The pre-training scenario may correspond to one of the original datasets or to the mixed unlabeled pool, whereas fine-tuning and evaluation are carried out on EPA500, EPA790, or ETU790. This protocol enables the analysis of both intra-scenario learning and cross-scenario transfer, which are discussed in Section 5.

5. Evaluation

5.1. Evaluation Protocol and Metrics

The evaluation considers the same three scenarios introduced in Section 3, namely EPA500, EPA790, and ETU790. For each experiment, the downstream classifier is fine-tuned on one labeled scenario and evaluated either on the same scenario or on a different one. Accordingly, two evaluation settings are considered:

- *Intra-scenario evaluation*: the fine-tuning and evaluation datasets coincide, i.e., $s_{ft} = s_{ev}$.
- *Cross-scenario evaluation*: the fine-tuning and evaluation datasets differ, i.e., $s_{ft} \neq s_{ev}$.

The SSL models are additionally characterized by the pre-training scenario $s_{pt} \in \{\mathrm{EPA500}, \mathrm{EPA790}, \mathrm{ETU790}, \mathrm{MIXED}\}$. The final reported results correspond to the configuration with six masked positions, selected from preliminary experiments. Two SSL objectives are compared: random token masking and contiguous masked-span prediction.

Performance is assessed through recall, specificity, and balanced accuracy (BAC). Let TP, TN, FP, and FN denote the true-positive, true-negative, false-positive, and false-negative counts, respectively. Then,

$$\mathrm{Recall} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}, \tag{8}$$

$$\mathrm{Specificity} = \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FP}}, \tag{9}$$

and

$$\mathrm{BAC} = \frac{\mathrm{Recall} + \mathrm{Specificity}}{2} \times 100. \tag{10}$$

Table 1. Intra-scenario results. The table reports the supervised baseline and the best SSL configuration for each evaluation scenario.

Scenario	Baseline			Best SSL					ΔBAC
	BAC (%)	Recall	Spec.	Objective	Pre-train	BAC (%)	Recall	Spec.	
EPA500	93.19	0.9715	0.8923	Contiguous masked-span	EPA500	97.27	0.9725	0.9730	+4.08
EPA790	96.78	0.9405	0.9951	Random masking	EPA790	97.01	0.9593	0.9809	+0.23
ETU790	96.77	0.9395	0.9960	Random masking	ETU790	97.32	0.9662	0.9802	+0.55

In the experiment logs, BAC is stored as weighted accuracy. Since the task is binary and the goal is to balance detection capability and false-alarm control, BAC is adopted as the main comparison metric throughout this section.

5.2. Same-Scenario Results

Table 1 summarizes the intra-scenario results. The supervised MLP baseline achieves BAC values of 93.19%, 96.78%, and 96.77% on EPA500, EPA790, and ETU790, respectively. These results indicate that the baseline is already strong in matched conditions, particularly on EPA790 and ETU790.

Nevertheless, the SSL-based models improve the best intra-scenario BAC in all three cases. On EPA500, the best result is obtained by contiguous masked-span prediction with EPA500 pre-training, reaching 97.27% BAC, with recall equal to 0.9725 and specificity equal to 0.9730. This corresponds to a gain of 4.08 percentage points over the baseline. On EPA790, the best result is achieved by random masking with EPA790 pre-training, yielding 97.01% BAC, 0.9593 recall, and 0.9809 specificity, which improves the baseline by 0.23 percentage points. On ETU790, the best configuration is also random masking, now with ETU790 pre-training, reaching 97.32% BAC, 0.9662 recall, and 0.9802 specificity, corresponding to a gain of 0.55 percentage points.

These results show that SSL is beneficial even in matched-domain conditions, although the magnitude of the gain depends on the scenario. The largest improvement occurs on EPA500, where the baseline exhibits comparatively lower specificity. By contrast, the gains on EPA790 and ETU790 are smaller, which suggests that the supervised baseline already captures much of the discriminative structure in those two scenarios. In these latter cases, the BAC improvement is mainly driven by higher recall, even when specificity decreases slightly.

5.3. Cross-Scenario Results

The cross-scenario results are more informative for assessing representation transfer. Table 2 reports the corresponding BAC, recall, and specificity values. In this setting, the average BAC of the supervised baseline across all transfer pairs is 92.72%, which confirms that generalization degrades when the fine-tuning and evaluation domains differ.

The strongest cross-scenario improvements are obtained with SSL pre-training. When the downstream classifier is fine-tuned on EPA500 and evaluated on EPA790, the best result is produced by random masking with MIXED pre-training, reaching 96.46% BAC, compared with 94.74% for the baseline. For the EPA500→ETU790 transfer, the same SSL configuration achieves 95.57% BAC, again improving over the baseline, which attains 93.54%.

Table 2. Cross-scenario results. For each transfer pair, the table reports the supervised baseline and the best SSL configuration.

Transfer pair	Baseline BAC (%)	SSL objective	Pre-train	SSL BAC (%)	Recall	Spec.	ΔBAC
EPA500 → EPA790	94.74	Random masking	MIXED	96.46	0.9512	0.9780	+1.72
EPA500 → ETU790	93.54	Random masking	MIXED	95.57	0.9148	0.9967	+2.04
EPA790 → EPA500	96.82	Random masking	EPA790	96.33	0.9552	0.9713	-0.49
EPA790 → ETU790	91.91	Random masking	ETU790	96.55	0.9376	0.9933	+4.64
ETU790 → EPA500	89.39	Contiguous masked-span	EPA790	95.96	0.9454	0.9738	+6.57
ETU790 → EPA790	89.93	Contiguous masked-span	EPA790	96.31	0.9480	0.9783	+6.38

For the EPA790→ETU790 transfer, the best result is obtained by random masking with ETU790 pre-training, yielding 96.55% BAC, compared with 91.91% for the baseline. This corresponds to a gain of 4.64 percentage points. The largest gains appear when the downstream model is fine-tuned on ETU790 and evaluated on the EPA scenarios. In that case, contiguous masked-span prediction with EPA790 pre-training reaches 95.96% BAC on ETU790→EPA500 and 96.31% BAC on ETU790→EPA790, whereas the baseline attains only 89.39% and 89.93%, respectively. These gains, equal to 6.57 and 6.38 percentage points, show that SSL substantially improves transfer under stronger distribution shifts.

The only transfer pair in which SSL does not surpass the baseline is EPA790→EPA500. In this case, the baseline itself achieves the best BAC, equal to 96.82%, while the best SSL result reaches 96.33%. This exception indicates that not all transfer directions benefit equally from self-supervised pre-training, and that some scenario pairs are already well aligned from the viewpoint of the downstream classifier.

Overall, the cross-scenario results show that SSL is particularly valuable under domain mismatch. This behavior is consistent with the main motivation of the paper: when labels are scarce and deployment conditions vary, the benefit of SSL lies not only in reducing annotation dependence, but also in learning representations that transfer more robustly across scenarios.

5.4. Effect of the Pre-Training Dataset and Masking Strategy

Additional trends emerge when the results are grouped by SSL configuration. For random masking, the MIXED pre-training set yields the best average cross-scenario BAC, reaching 95.35%. This suggests that exposing the encoder to a more heterogeneous unlabeled pool improves robustness under transfer. For contiguous masked-span prediction, the best average cross-scenario BAC is obtained with EPA790 pre-training, reaching 94.15%. In contrast, ETU790-only pre-training produces the weakest average transfer performance for both SSL objectives.

Another relevant observation is that no single masking strategy dominates all scenarios. Contiguous masked-span prediction provides the best results in several transfer settings, especially when the downstream model is fine-tuned on ETU790 and evaluated on EPA500 or EPA790. However, random masking combined with MIXED pre-training provides the strongest overall average transfer performance and also yields the best results for EPA500-based transfer. Therefore, the results do not support a universal conclusion in favor of one masking strategy alone. Instead, they indicate that the effectiveness of the SSL objective depends on the interaction among pre-training, fine-tuning, and evaluation domains.

5.5. Discussion

The evaluation supports three main conclusions. First, SSL improves the downstream classifier in all intra-scenario evaluations and in five of the six cross-scenario transfer pairs, even under severe label scarcity. Second, the gains are modest in matched conditions but become substantially larger in cross-scenario transfer, which is precisely the setting in which annotation scarcity and domain shift are most problematic. Third, the composition of the unlabeled pre-training set is a key factor: heterogeneous pre-training through MIXED data is highly beneficial for random masking, while EPA790 emerges as the most transferable single-scenario source for contiguous masked-span prediction.

These findings suggest that the main advantage of SSL in this problem is not limited to replacing labels during pre-training. Rather, SSL reshapes the PDP representation space in a way that improves robustness across propagation conditions and deployment configurations. This property is particularly relevant for early collision detection in practical CIoT systems, where collecting large unlabeled signal traces is considerably easier than generating new labeled datasets through detailed RA simulation and Msg3-based collision confirmation.

Regarding the computational cost, while SSL introduces a pre-training stage, this process is performed entirely on unlabeled data, which is abundant. The overhead is compensated by the fact that the downstream fine-tuning converges $5\times$ faster than the baseline trained from scratch. The superiority of contiguous masked-span prediction in high-complexity channels (ETU790) suggests that local structural correlations in the PDP are more indicative of collisions in multi-path environments. Conversely, random masking benefits from the MIXED dataset because it forces the model to learn more global, scenario-agnostic features, which explains its better average performance across diverse transfer pairs.

6. Conclusions

This paper investigated self-supervised learning for early PRACH preamble collision detection under realistic CIoT random access conditions. The study was motivated by an important practical limitation of supervised approaches: collision labels are costly to obtain because they depend on detailed simulation and protocol-level verification beyond Msg1 reception. To address this issue, a two-stage pipeline was adopted, in which a Transformer-based encoder is first pre-trained on unlabeled PDP data and then representations are transferred to a downstream classification task fine-tuned with a small labeled data subset.

The evaluation considered the three scenarios EPA500, EPA790, and ETU790, under both same-scenario and cross-scenario settings. The results showed that self-supervised pre-training improves the downstream classifier in all same-scenario evaluations and in five of the six cross-scenario transfer pairs. These gains were modest when fine-tuning and evaluation were performed on the same scenario but became more pronounced under cross-scenario transfer, where domain mismatch is stronger and labeled data are more difficult to obtain. This behavior indicates that the main advantage of SSL in this problem is not limited to reducing dependence on labeled data but also includes improving the robustness of the learned representation across different propagation conditions.

The experiments also showed that the effectiveness of SSL depends on both the masking strategy and the pre-training domain. Contiguous masked-span prediction yielded the best results in several transfer configurations, particularly when the downstream model was fine-tuned on ETU790 and evaluated on EPA500 or EPA790. In contrast, random token masking combined with MIXED pre-training achieved the strongest average cross-scenario performance, suggesting that exposure to a more heterogeneous unlabeled pool improves representation transfer, confirming the practical value of masking-based pre-training for this task.

Overall, the findings support the use of SSL as a viable alternative to purely supervised learning for early preamble collision detection. By exploiting unlabeled PDP observations during pre-training, the proposed pipeline reduces reliance on expensive labels while improving downstream performance in most evaluated conditions. This property is particularly relevant for practical CIoT deployments, where collecting large volumes of raw signal measurements is considerably easier than generating new annotated datasets for each scenario.

As future work, evaluate model compression strategies such as quantization and pruning, and investigate additional self-supervised pretext tasks, including contrastive, destillation, hybrid, and alternative masking-based formulations. Another important direction is to assess the proposed approach under a broader range of channel models and deployment conditions, including vehicular networks and 5G LEO satellite scenarios, in which mobility and varying channel conditions make robustness and transferability relevant.

Acknowledgements

This work was supported by the INCT of Intelligent Communications Networks and the Internet of Things (ICoNIoT) via CNPq (405940/2022-0) and CAPES (88887.954253/2024-00); the FAPESP Tematic Project IDRIC (Distributed Intelligence in Communication Networks and the Internet of Things) process 23/00673-7; UNICAMP's PIND (3406/23) and BTG Pactual.

References

3GPP (2011). Study on ran improvements for machine-type communications. Technical Report TR 37.868 V11.0.0, 3rd Generation Partnership Project (3GPP).

Cardenas, G. M., Gonzalez, D. C., Guevara, J. C., Astudillo, C. A., and da Fonseca, N. L. S. (2025). Ml-based preamble collision detection in the random access procedure of cellular IoT networks. *arXiv preprint arXiv:2510.25145*.

de Andrade, T. P., Astudillo, C. A., Sekijima, L. R., and da Fonseca, N. L. S. (2017). The random access procedure in long term evolution networks for the internet of things. *IEEE Communications Magazine*, 55(3):124–131.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational*

Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186.

Dong, J., Wu, H., Zhang, H., Zhang, L., Wang, J., and Long, M. (2023). SimMTM: A simple pre-training framework for masked time-series modeling. In *Advances in Neural Information Processing Systems*, volume 36.

Enoki, G. S. and Astudillo, C. A. (2025). Self-supervised learning for spectrogram-based identification of mobile communications technologies. In *Wireless Days Conference (WD)*.

Hajiramezanali, E., Diamant, N., Scalia, G., and Shen, M. W. (2023). STab: Self-supervised learning for tabular data. OpenReview. Available at OpenReview.

Jang, H. S., Lee, H., Quek, T. Q. S., and Shin, H. (2021). Deep learning-based cellular random access framework. *IEEE Transactions on Wireless Communications*, 20(11):7503–7518.

Li, X., Zhao, Z., Zhang, Y., Zheng, S., and Dai, S. (2023a). Spectrum sensing algorithm based on self-supervised contrast learning. *Electronics*, 12(6):1317.

Li, Z., Rao, Z., Pan, L., Wang, P., and Xu, Z. (2023b). Ti-MAE: Self-supervised masked time series autoencoders. *arXiv preprint arXiv:2301.08871*.

Magrin, D., Pielli, C., Stefanovic, C., and Zorzi, M. (2019). Enabling LTE RACH collision multiplicity detection via machine learning. In *2019 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, pages 1–8.

Milosheski, L., Cerar, G., Bertalanič, B., Fortuna, C., and Mohorčič, M. (2023). Self-supervised learning for clustering of wireless spectrum activity. *Computer Communications*, 212:353–365.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.

Wei, C., Bianchi, G., and Cheng, R. (2015). Modeling and analysis of random access channels with bursty arrivals in OFDMA wireless networks. *IEEE Transactions on Wireless Communications*, 14(4):1940–1953.

Yin, Y., Zhao, D., Li, X., and Zeng, S. (2023). Learning based preamble collision detection of cellular random access by physical layer features. In *2023 International Conference on Networking and Network Applications (NaNA)*, pages 28–33.