Research Article

ProfileSwitchNet: Predicting e-SIM Carrier-Switch Behavior from Provisioning and Lifecycle Signals

Aqib Hassan [1*], and Mubashir Gulab [2]

[1] Director At Future Syncs Limited UK
[2] School of Law and Social Sciences, University of East London
[*]Corresponding Author: Aqib Hassan Author Email: Hassan@Futuresyncs.com
Received: March 14, 2025 **Accepted:** December 17, 2025 **Published:** December 25, 2025

Abstract: This work introduces ProfileSwitchNet, a predictive model for forecasting e-SIM carrier-switch behavior based on historical provisioning and lifecycle data. Each e-SIM profile is represented by features such as device type, initial and subsequent carrier profiles, activation counts, roaming flags, region, data plan class, and recorded status changes over time. A sequence-aware model (e.g., temporal gradient boosting or a recurrent neural network) is trained to estimate the probability that a profile will be deactivated or switched to a different carrier within a future time window. The framework evaluates performance using ROC-AUC, precision–recall, and calibration metrics, and analyzes feature importance to identify the strongest drivers of switching, such as device category or activation failure bursts. ProfileSwitchNet is intended to help operators design better retention offers, optimize roaming strategies, and understand early signals of churn in eSIM ecosystems.

***Keywords*:** ProfileSwitchNet, eSIM carrier-switch, ROC-AUC, precision–recall, calibration metrics, activation failure bursts, churn

1. Introduction

The adoption of embedded SIM (eSIM) and remote SIM provisioning (RSP) has revolutionized the manner in which mobile subscriptions are provided and managed through smart phones, IoT systems and M2M implementations. Recent security considerations of the consumer RSP protocol have pointed to its flexibility and its complexity, and additional areas of attack have been brought into the limelight of profile download, activation, and profile management processes [1]. Additional studies on transparency and privacy of eSIM profiles demonstrate that the communication between the device, subscription manager, and mobile network operator is stipulated by complex trust parameters and protocol states that cannot necessarily be seen by users or operators [2].

Simultaneously, eSIM is now an essential facilitator of secure connections in 5G and IoT applications. Suggested eSIM-based credentialing provisioning and authentication schemes of the IoT devices highlight the usefulness of remote profile lifecycle management but also observe the operational and security complications with large scale implementation [3]. Regional studies of eSIM adoption, including those in Latin America, highlight the benefits of the business in terms of enhanced roaming, multi-operator coverage, and more cost-efficient logistics, with barriers including regulation, interoperability, and ecosystem maturity also being recorded [4]. Long risk-wise, systematic reviews based on ISO 31000 have emphasized that eSIM technology may introduce new operational risks in the areas of failure in provisioning, misconfigurations and reliance on remote lifecycle operations, which have to be properly identified and addressed [5].

Although this paper is characterized by increased research on the subject of eSIM technology, security, and operational risk, there was a lack of attention as far as data-driven modelling of eSIM carrier-switch behavior is concerned. Most available literature takes eSIM either in a protocol-security, or qualitative risk-management perspective. Meanwhile, customer churn prediction and behavioral modelling on a subscriber basis have been the subject of ex-post facto discussion in the telecom analytics literature, typically based on billing, usage, and complaint data, including none of the finer-grained eSIM provisioning and lifecycle logs. This introduces a methodological and practical disjuncture: operators do not have predictive tools that leverage the distinctive cues that exist in eSIM lifecycle events (e.g., counts of profile activations, changes in a profile as a result of roaming, recurring profile changing errors, etc.) to pre-empt carrier switching and early churn in an eSIM ecosystem.

In response to this gap, we have suggested ProfileSwitchNet, a predictive architecture that uses eSIM provisioning and lifecycle cues to predict the likelihood of deactivation or switching a profile to a new carrier in the next time window. Each eSIM profile is modelled as a time series of lifecycle events and contextual attributes (type of device, first carrier, second carrier, roaming status, region, data plan type, activation failures), allowing sequence-aware models to model dynamic behavior instead of snapshot models. This approach of specifically targeting carrier-switch events rather than churn at the subscriber level, ProfileSwitchNet is addressing a growing operational issue specific to the operation of networks made possible by eSIM.

1.1 Contributions

The main contributions of this paper are as follows:

- **eSIM-centric problem formulation for carrier switching.** We formalize the problem of predicting eSIM carrier-switch behavior using provisioning and lifecycle logs, bridging a gap between eSIM/RSP research which has focused on protocol security, risk, and deployment issues [1]–[5] and telecom churn prediction, which largely operates on traditional billing and usage data.
- **Sequence-aware representation of eSIM lifecycle signals.** We design a feature representation in which each eSIM profile is modelled as a **temporal sequence** of lifecycle events (downloads, activations, deactivations, roaming-related modifications, and error events), enriched with device and plan attributes. This allows temporal gradient boosting or recurrent neural networks to learn switching patterns that static models cannot capture.
- **Predictive framework and evaluation for carrier-switch risk.** We develop ProfileSwitchNet, a sequence-aware predictive framework that estimates the probability of profile deactivation or carrier switch within a specified horizon. The framework evaluates performance using ROC-AUC, precision–recall, and calibration metrics, and employs feature importance and temporal attribution analysis to identify the most salient drivers of switching (e.g., device category, activation failure bursts, roaming patterns).
- **Bridging churn analytics and eSIM lifecycle management**. By adapting ideas from telecom churn modelling and survival analysis to eSIM lifecycle data, we demonstrate how carrier-switch risk can be anticipated earlier and at a finer granularity than traditional subscriber-level churn. This supports more targeted retention offers, optimized roaming strategies, and proactive monitoring of operational risk in eSIM ecosystems.

The remainder of this paper is organized as follows. Section 2 reviews related work on eSIM technology and remote provisioning, telecom churn prediction using machine learning and deep learning, and time-to-churn and switching behavior modelling. Section 3 presents the ProfileSwitchNet problem formulation and model architecture. Section 4 describes the experimental setup, including datasets, feature engineering, baselines, and evaluation metrics. Section 5 reports and discusses the empirical results, and Section 6 concludes the paper and outlines directions for future work.

2. Literature Review

This section reviews prior work in three main areas: (i) eSIM and remote SIM provisioning technology and operational risk, (ii) machine learning and deep learning for telecom churn prediction, and (iii) modelling of switching and time-to-churn behavior. We then summarize the key gaps with respect to eSIM carrier-switch prediction and position ProfileSwitchNet within this landscape.

2.1 eSIM Technology, Remote Provisioning, and Risk

Security analyses of consumer RSP protocols have exposed subtle vulnerabilities and trust assumptions within the eSIM ecosystem, emphasizing the need for rigorous design and verification of profile download, installation, and management procedures [1]. Complementary work on SIM profile transparency has shown that subscription managers and operators must balance privacy, accountability, and operational control, with protocol details directly influencing what information is visible and auditable [2].

In the context of IoT and 5G, eSIM has been proposed as a secure anchor for remote credential provisioning and device authentication. Protocols that leverage eSIM for IoT connectivity aim to improve scalability and security, but they also highlight challenges around large-scale lifecycle management and interoperability across multiple operators [3]. Regional deployment studies, such as those focusing on Latin America, provide empirical evidence of how eSIM adoption is shaped by regulatory constraints, device availability, and operator readiness, as well as the perceived advantages of simplified roaming and multi-carrier support [4].

From a risk-management perspective, systematic reviews based on ISO 31000 indicate that eSIM introduces new operational and strategic risks ranging from provisioning failures and incorrect profile management to dependence on third-party subscription managers that must be identified, assessed, and mitigated as part of an integrated risk framework [5]. However, these works remain largely conceptual or protocol-oriented, without exploiting eSIM provisioning and lifecycle logs as input to predictive models that anticipate carrier-switch or churn behavior. ProfileSwitchNet builds on this foundation by treating provisioning and lifecycle events as a rich source of behavioral signals for data-driven risk and churn analytics.

2.2 Machine Learning for Telecom Churn Prediction

Telecom churn prediction is a mature area where machine learning has been extensively applied to subscriber-level data. Early work demonstrated that data mining techniques can significantly improve churn prediction by exploiting billing, usage, and demographic features [6]. Subsequent studies have developed improved models and frameworks, using advanced preprocessing, feature engineering, and ensemble methods to enhance predictive performance [7], [8].

With the growth of big data platforms, researchers began to explore large-scale, distributed churn prediction pipelines. For example, big data architectures have been used to train and deploy machine learning models over high-volume call and usage datasets, enabling near-real-time churn analytics [9]. More recent contributions focus on combining traditional models with modern optimization and feature-selection strategies to further improve performance and robustness in practical settings [10], [11], [26].

Deep learning has also been adopted for telecom churn prediction. Studies have proposed deep architectures that integrate behavioral features, temporal patterns, and customer profiles to capture complex nonlinear relationships [12], [13]. Hybrid and hierarchical deep learning approaches further combine multiple feature types or embed ensemble strategies within deep architectures to leverage complementary information sources [14]. Other work explores specialized recurrent models, such as Swish RNN-based architectures, to better encode sequential usage patterns [15].

Ensemble-based approaches remain prominent, where clustering and classification are combined to segment customers and tailor churn models to each segment [16]. Integrated frameworks that unify churn prediction and customer segmentation have been proposed to support more actionable marketing and retention strategies within telecom businesses [17]. Comprehensive comparative studies evaluating large numbers of classifiers over telecom churn datasets—help establish strong baselines and guide model selection for practical deployments [18].

More recently, there has been a growing emphasis on explainability and business interpretability in churn modelling. Approaches that incorporate SHAP values or other explanation techniques provide insight into the main drivers of churn and help align model predictions with domain expertise [19]. Complementary work proposes end-to-end customer churn prediction systems tailored to business constraints, emphasizing model integration, performance, and maintainability within operational environments [20].

While this body of work demonstrates the effectiveness of machine learning and deep learning for subscriber churn prediction, it typically relies on high-level billing and usage data, with limited access to low-level SIM or eSIM lifecycle events. ProfileSwitchNet draws inspiration from these methods but shifts focus to eSIM profile-level switching behavior, leveraging provisioning and lifecycle signals that are unique to eSIM-based architectures.

2.3 Switching and Time-to-Churn Behavior

Beyond binary churn prediction, several studies have examined the determinants and temporal dynamics of customer switching. Early research in mobile telephony markets analyzed how service quality, pricing, and customer loyalty influence subscriber churn, providing empirical evidence that both technical and relational factors drive switching behavior [21].

Based on this, machine learning has been integrated with push-pull-mooring models to study switching among customers in telecom markets, emphasizing interactions among dissatisfaction, attractiveness of substitutes, and personal or contextual moorings [22].

Another significant body of work is time-to-churn modelling. Other research works have adopted the social network analysis and survival modelling to forecast the time to churn in prepaid mobile markets, resulting in results that the network position and peer behavior significantly influence the churn risk [23]. A variation of this has also been performed to model time-to-churn directly using cox regression and related survival analysis methods to allow more finer risk estimation over time instead of fixed churn labels [24].

These works do not deal directly with eSIM, but they emphasize the benefit of time and behavioral modelling in the behavior of switching and churn. Other digital sectors, including online gambling, have applied, on the churn side, related research with sophisticated machine-learning methods to forecast and determine high-churn users in complex transactional settings [25]. These studies, however, do not focus specifically on the events of eSIM carrier-switch based on provisioning and lifecycle logs or model specifically the operational context of a eSIM-based subscription management.

2.4 Summary and Research Gap

Across the reviewed literature, three key observations emerge:

1. **eSIM and RSP research is protocol and risk-centric.** Prior work focuses on security analysis, transparency, IoT credential provisioning, regional adoption, and risk management for eSIM [1]–[5]. These studies do not develop predictive models that use eSIM lifecycle data to anticipate carrier-switch or churn behavior.
2. **Telecom churn models operate at subscriber level using traditional data.** Machine-learning and deep-learning approaches for churn prediction [6]–[20] primarily exploit billing, usage, complaint, and demographic features. They do not leverage eSIM provisioning or lifecycle logs, and they typically predict subscriber churn rather than profile-level carrier switching in multi-carrier eSIM settings.
3. **Temporal and behavioral models do not use eSIM lifecycle signals.** Time-to-churn, switching behavior, and survival analysis [21]–[25] demonstrate the importance of temporal dynamics and network context, but they have not been applied to eSIM-specific lifecycle data or cross-carrier switching events.

ProfileSwitchNet addresses these gaps by (i) framing eSIM carrier-switch prediction as a sequence modelling problem over provisioning and lifecycle signals, (ii) adapting churn and survival modelling concepts to profile-level eSIM data, and (iii) providing interpretable importance analyses for operational drivers of switching (e.g., activation failure bursts, roaming patterns, device categories).

Table 1. Summary of related work and research gaps with respect to eSIM carrier-switch prediction

Category / Representative work	Main focus	Key limitations for eSIM carrier-switch prediction	How ProfileSwitchNet addresses the gap
eSIM security, transparency, IoT provisioning, and risk [1]–[5]	Protocol security of RSP, transparency of SIM profiles, eSIM for IoT and 5G, regional adoption, ISO 31000-based operational risk analysis.	Focus on protocol design, qualitative risk assessment, and deployment issues; do not build predictive models over provisioning/lifecycle	Uses eSIM provisioning and lifecycle logs as primary features and formulates carrier-switch prediction as a data-driven sequence modelling problem at profile level.

		logs; do not target carrier-switch behavior.	
Telecom churn prediction using ML and DL [6]–[20]	Subscriber-level churn prediction using billing, usage, demographic and sometimes complaint data; includes deep, ensemble, and explainable ML frameworks for telco churn.	Operate at subscriber granularity; rely on traditional data sources (usage/billing); do not exploit eSIM lifecycle signals; typically predict churn vs. non-churn, not which carrier a profile will switch from/to in an eSIM ecosystem.	Adapts churn modelling ideas to eSIM profiles, incorporating provisioning events, carrier changes, roaming flags, and activation failures; predicts profile-level switch/deactivation risk over a future horizon.
Behavioral switching and time-to-churn modelling [21]–[25]	Determinants of switching in mobile markets; push–pull–mooring frameworks; survival analysis and time-to-churn using network and behavioral features; advanced churn modelling in other digital domains.	Highlight importance of temporal dynamics and context but do not use eSIM lifecycle data; do not model carrier-switch events specific to eSIM RSP; typically operate on aggregated customer behavior rather than low-level lifecycle logs.	Incorporates temporal and behavioral modelling (sequence-aware models, risk over future window) directly over eSIM lifecycle sequences, enabling early detection of profiles likely to switch carrier or be deactivated.

This gap analysis motivates the design of ProfileSwitchNet as an eSIM-first, sequence-aware predictive framework that integrates ideas from telecom churn modelling, temporal risk analysis, and eSIM lifecycle management to anticipate carrier-switch behavior in modern eSIM ecosystems.

3. Methodology

This section presents the ProfileSwitchNet methodology for predicting eSIM carrier-switch behavior from provisioning and lifecycle logs. We first formalize the prediction problem, then describe the dataset and preprocessing pipeline, followed by feature engineering, model architecture, and training/evaluation protocol.

3.1 Problem Formulation

Let $\mathcal{P}$ denote the set of eSIM profiles observed in the provisioning logs. Each $p \in \mathcal{P}$ is associated with a time-ordered sequence of lifecycle events

$$\mathcal{E}_p = e_p 1, e_p 2, \dots \dots \dots, e_p, T_p$$

where each event e_p, t corresponds to a provisioning or lifecycle operation such as profile download, enable, disable, delete, or update, and is timestamped.

Each event is described by a feature vector $e_p, t \in \mathbb{R}^{d_e}$ encoding:

- **Event metadata** (event type, status, error code, method),
- **Carrier context** (current carrier, previous carrier, plan tier),
- **Device context** (device type, consumer vs IoT/M2M, region),

- **Temporal information** (inter-event time, lifecycle age).

In addition, each profile has a set of static attributes $s_p \in \mathbb{R}^{d_s}$ that are constant (or slowly changing) over the observation period, such as device form factor, initial carrier, and base region.

We assume an observation window $[0, T_{obs}]$ and a prediction horizon of length ΔT. For each profile ppp, we construct:

- An input sequence $\mathcal{E}_p^{obs}$containing all events with timestamps $\leq T_{obs}$.
- A binary label $\mathcal{Y}_\mathcal{P}$ defined as:
  - $\mathcal{Y}_\mathcal{P} = 1$ if a carrier-switch or profile deactivation event occurs in
  - $\mathcal{Y}_\mathcal{P} = 0$ otherwise (no switch/deactivation within the horizon).

Profiles that churn/switch before the observation window or lack sufficient history may be treated as filtered or censored, depending on the experimental design (see Section 3.2.3).
ProfileSwitchNet learns a function

$$f_0 : (\mathcal{E}_p^{obs}, s_p) \mapsto \widehat{\mathcal{Y}_\mathcal{P}} \in [0,1]$$

parameterized by θ, such that $\widehat{\mathcal{Y}_\mathcal{P}}$ estimates the probability that profile p will switch carriers or be deactivated within the next ΔT units of time. The model is trained to minimize a class-imbalance–aware cross-entropy or focal loss over all profiles in the training set.

3.2 Datasets and Preprocessing

3.2.1 eSIM/eUICC Remote Provisioning (RSP) Synthetic Dataset

We base our experiments on the eSIM/eUICC Remote Provisioning (RSP) – Synthetic Dataset, which contains 100% synthetic, privacy-compliant provisioning events for consumer, IoT, and M2M devices. According to the dataset description, it includes 100 records and 18 columns and is designed for analyzing eSIM provisioning methods, carrier switching, and lifecycle management in telecommunications scenarios.
The dataset provides, at event level, information such as:

- Unique identifiers for profiles and devices,
- Carrier or subscription identifiers,
- Provisioning operations (e.g., download, enable, disable, delete),
- Status indicators (success/failure),
- Timestamps and possibly region/country tags,
- Support for multi-profile management (multiple carriers per device).

Because profiles may appear in multiple rows (one per event), we aggregate event records into profile-centric sequences as described in Section 3.2.2. Table 2 summarizes the main characteristics of the dataset after aggregation to profile level.

Table 2. Summary of eSIM provisioning dataset .

Statistic	Value (example)
Number of profiles	80
Number of devices	60
Number of carriers	5
Total provisioning events	100
Mean events per profile	1.25
Median events per profile	1
Profiles with ≥ 2 carriers (multi-carrier)	20
Switch/deactivation positives (y=1y=1y=1)	24
Non-switch profiles (y=0y=0y=0)	56

If you augment the dataset (e.g., by generating additional synthetic sequences or integrating operator logs), you should update Table 2 accordingly.

3.2.2 From Events to Profile Sequences

To construct input sequences for ProfileSwitchNet, we perform the following steps:

1. Sorting and grouping.
    - Group rows by profile identifier (e.g., profile_id) and sort events by timestamp.
    - For each profile p , obtain a time-ordered event sequence $\mathcal{E}_p$.
2. Defining observation and horizon windows.
    - For static experiments, choose a fixed observation cutoff T_{obs} (e.g., last timestamp minus ΔT).
    - Keep all events with $t < \Delta T$ as input, and search for carrier-switch or deactivation events in $(T_{obs}, T_{obs} + \Delta T)$
3. Label construction.
    - Mark $\mathcal{Y}_{\mathcal{P}} = 1$ if any switch or deactivation event occurs in the horizon; otherwise, $\mathcal{Y}_{\mathcal{P}} = 0$
    - Optionally, distinguish carrier switch from deactivation with a multi-class label if you want a more fine-grained study.
4. Sequence truncation and padding.
    - Set a maximum sequence length L_{max} (e.g., 16 events).
    - For profiles with longer histories, keep the most recent L_{max} events; for shorter ones, apply left-padding with a special “no-event” token or mask.
5. Train/validation/test split.
    - Split by profile (not event) to avoid leakage, e.g., 60% train, 20% validation, 20% test.

- Optionally, maintain carrier or region stratification to preserve distributional balance.

3.2.3 Handling Censored Profiles and Class Imbalance

Some profiles may not have enough observation history (e.g., only one event right at the end of the log) or may churn before the observation window. These can be:

- Excluded as insufficient data, or
- Treated as right-censored cases if you adopt a survival-analysis variant of ProfileSwitchNet.

Class imbalance is common—carrier-switch events are typically rarer than non-switch profiles. We address this by:

- Using a class-weighted loss or focal loss during training,
- Optionally oversampling rare positives in the training set,
- Monitoring precision–recall and calibration metrics, not just accuracy.

3.3 Feature Engineering

Feature engineering in ProfileSwitchNet is organized into two main groups:

- Static profile-level features S_p,
- Event-level temporal features $x_{p,t}$.

3.3.1 Static Profile-Level Features

Static features describe relatively stable properties of the profile and associated device:

- Device & user segment:
  - Device type (smartphone, tablet, wearable, IoT sensor, vehicle modem),
  - Segment flag (consumer vs IoT/M2M).
- Initial subscription context:
  - Initial carrier identifier,
  - Initial plan class (prepaid/postpaid, data-only vs voice+data),
  - Region or country of first activation.
- Lifecycle summary statistics:
  - Total number of activation attempts,
  - Total number of failures,
  - Number of carriers ever used,
  - Time since first provisioning event.

Categorical variables (e.g., carrier, device type, region) are encoded using learned embeddings or one-hot encoding; continuous variables are standardized.

3.3.2 Event-Level Temporal Features

Event-level features represent each provisioning or lifecycle event:

- Event type: download, enable, disable, delete, update, error.
- Event outcome: success vs failure, error code category.
- Carrier & plan context: carrier involved in the event, plan tier, roaming flag.
- Temporal dynamics:
  - Time since previous event,
  - Cumulative lifecycle age at event time,
  - Time since last carrier switch (if any).

- Operational indicators:
    - Number of consecutive failures up to time t,
    - Flag for "late night" / "peak time" provisioning.

These features are combined into an event vector $x_{p,t}$ Categorical fields are embedded; continuous fields are normalized.

To keep the description structured, **Table 3** summarizes the main feature groups.

Table 3. Feature categories and examples used in ProfileSwitchNet.

Feature group	Level	Example features
Device & segment	Static	Device type, OS category, consumer vs IoT/M2M flag
Initial subscription	Static	Initial carrier, initial plan class, initial region/country
Lifecycle summary	Static	Total activations, total failures, number of carriers used, lifetime (days)
Event type & status	Temporal	Event type (download/enable/disable/delete/update), success/failure flag, error code category
Carrier & plan context	Temporal	Carrier at event, plan tier, roaming flag, multi-profile indicator
Temporal dynamics	Temporal	Time since previous event, lifecycle age, time since last switch
Operational burst signals	Temporal	Rolling count of failures in last kkk events, rapid profile flips, off-hours provisioning

By combining these features across time, ProfileSwitchNet can learn patterns such as "frequent failures followed by cross-border roaming" as precursors to carrier switching.

3.4 ProfileSwitchNet Architecture

ProfileSwitchNet is designed as a dual-branch sequence-aware architecture that fuses static context and temporal behavior.

3.4.1 Overview of the Pipeline

At a high level, the pipeline consists of:

1. Log ingestion & preprocessing (Section 3.2)
2. Feature extraction (static and temporal) (Section 3.3)
3. Dual-branch model:
    - Static context encoder
    - Temporal sequence encoder
4. Fusion & prediction of carrier-switch risk.

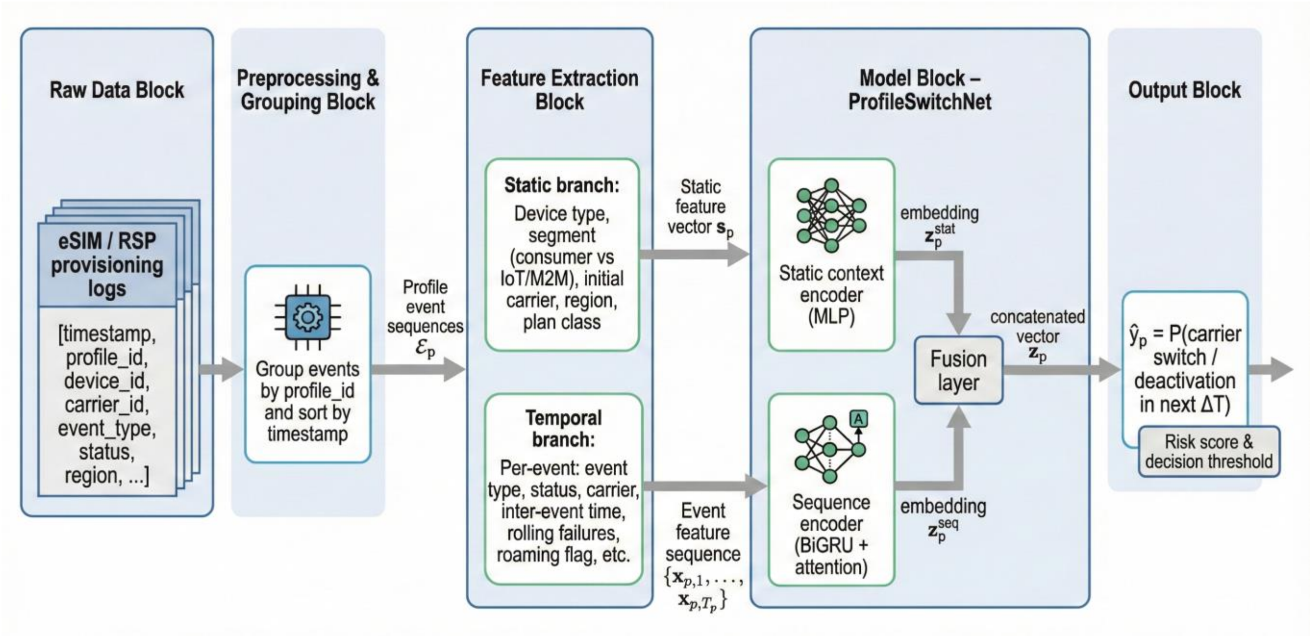

Figure 1. High-level ProfileSwitchNet pipeline (placeholder).

Figure 1 should show a flow from raw provisioning logs → event grouping by profile → feature extraction (static + event-level) → dual-branch model (static encoder + temporal encoder) → fusion layer → predicted probability of carrier switch/deactivation.

3.4.2 Temporal Sequence Encoder

For each profile $\mathcal{P}$, the event sequence $\mathcal{X}_p.1, \ldots\ldots\ldots..\mathcal{X}_p, \mathcal{T}_p$ is passed through:

1. **Embedding and projection.**
   - Categorical fields (event type, carrier, etc.) are mapped to embeddings.
   - Continuous fields are concatenated and linearly projected.
   - The final per-event input dimension is d_{in}.
2. **Recurrent or temporal model.**
   - A **bidirectional GRU** (BiGRU) or LSTM processes the event sequence:

$$h_{p,t} = BiGRU(\mathcal{X}_{p,t}, h_{p,t-1})$$

   - This captures forward and backward temporal dependencies across events.
3. **Attention-based aggregation.**
   - An attention mechanism assigns a weight $\alpha_{p,t}$ to each time step:

$$\alpha_{p,t} = \frac{\exp\left(u^{\mathrm{T}} \tanh\left(W_\alpha h_{p,t}\right)\right)}{\sum_{j=1}^{T_p} exp\left(u^T tanh\left(W_\alpha h_{p,j}\right)\right)}$$

   - The sequence embedding is:

$$\mathbf{z}_p^{\mathrm{seq}} = \sum_{t=1}^{T_p} \alpha_{p,t} \mathbf{h}_{p,t}$$

This sequence embedding encodes both **what** happened (event types, carriers, errors) and **when/how often** (temporal dynamics).

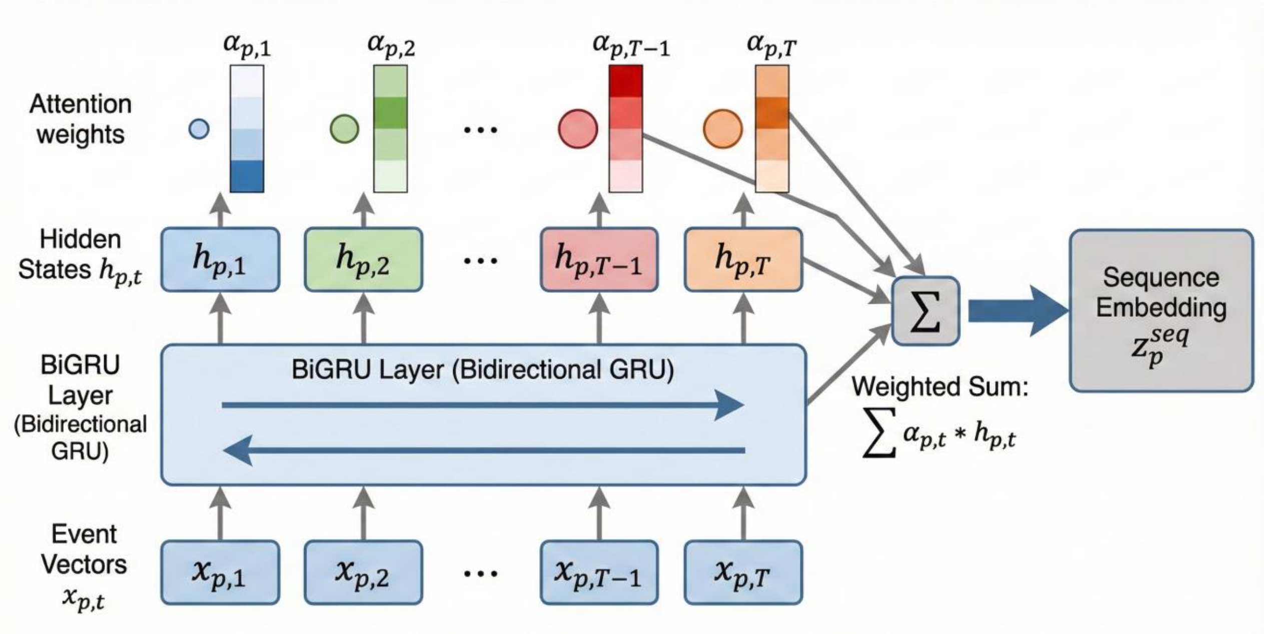

Figure 2. Temporal sequence encoder of ProfileSwitchNet.

Figure 2 should depict event vectors flowing into a BiGRU layer, followed by an attention layer that produces a single sequence embedding $\mathbf{z}_p^{\text{seq}}$. A heatmap overlay on the timeline can represent attention weights, showing which events the model focuses on for switch prediction.

3.4.3 Static Context Encoder

Static features S_p are passed through a small feed-forward network:

$$\mathbf{z}_p^{\text{stat}} = \sigma\big(W_2 \phi\big(W_1 \mathbf{s}_p + \mathbf{b}_1\big) + \mathbf{b}_2\big)$$

where ϕ is a non-linear activation (e.g., ReLU), σ can be ReLU or GELU, and dropout is applied between layers for regularization. This branch encodes the **background context** (device type, initial carrier, initial region, etc.) that modulates switching behavior.

3.4.4 Fusion and Output Layer

The two embeddings are concatenated:

$$[\mathrm{z}_p = [\mathrm{z}_p^{\text{seq}} \,||\, \mathrm{z}_p^{\text{stat}}\,]$$

and fed into a prediction head:

$$\widehat{y_{\mathcal{P}}} = \sigma\big(W_o \mathbf{z}_p + \mathbf{b}_o\big)$$

where σ is the sigmoid function. The model is trained with a class-weighted binary cross-entropy:

$$\mathcal{L} = -\sum_{p} w_{y_p}\left(y_p \log\, \hat{y}_p + \left(1 - y_p\right)\log\left(1 - \hat{y}_p\right)\right)$$

where $w_1 > w_0$ to account for the relative rarity of switch events.

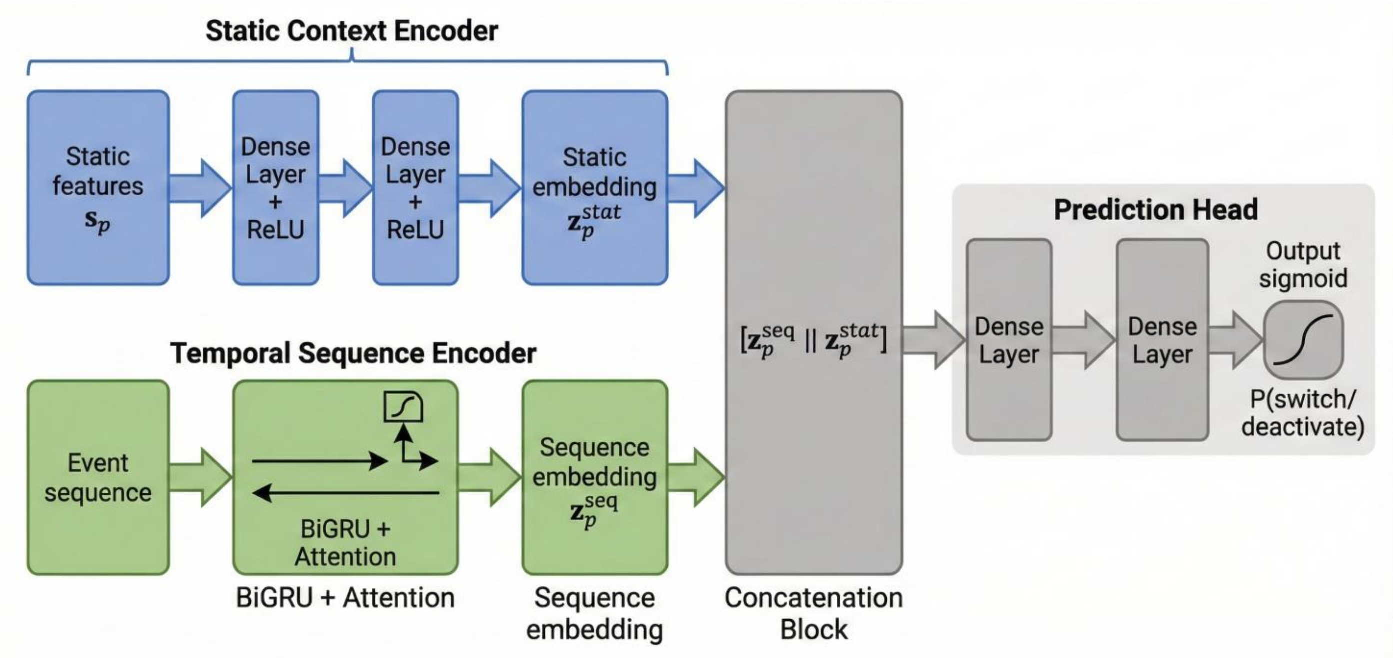

Figure 3. Dual-branch architecture of ProfileSwitchNet

Figure 3 should show static features entering a small MLP branch, temporal events entering the BiGRU+attention branch, and both branches feeding into a fusion layer and a final sigmoid output node representing switch probability.

3.5 Training Configuration and Baselines

We benchmark ProfileSwitchNet against simpler models that operate on **aggregated static features only**, providing insight into the added value of temporal modelling.

3.5.1 Baseline Models

- **Logistic Regression (LR).**
  Uses static and aggregated temporal features (e.g., total failures, number of carriers, last carrier, last event type) without sequence modelling.
- **Gradient Boosted Trees (e.g., XGBoost/LightGBM).**
  Operates on the same static and aggregated features, capturing nonlinear interactions.
- **Shallow MLP.**
  A multi-layer perceptron trained on flattened features with no explicit temporal modelling.

ProfileSwitchNet extends these baselines by modelling **full event sequences** and learning attention weights over them.

3.5.2 Hyperparameters and Training

Hyperparameters are tuned on the validation set using grid or Bayesian search. An example configuration is shown in Table 4.

Table 4. Example hyperparameters for ProfileSwitchNet and baselines.

Model	Key hyperparameters (example)
Logistic Regression	Penalty = L2; C = 1.0; class_weight = "balanced"; solver = "lbfgs"
Gradient Boosting	Trees = 200; max_depth = 4; learning_rate = 0.05; subsample = 0.8
MLP (baseline)	Hidden layers = [64, 32]; activation = ReLU; dropout = 0.3; batch_size = 64
ProfileSwitchNet	Event embedding dim = 32; BiGRU hidden size = 64 (bidirectional); static branch hidden = 32; dropout = 0.3; optimizer = Adam (lr = 1e-3); batch_size = 32; epochs = 50 with early stopping

3.5.3 Evaluation Metrics

We evaluate all models using:

- ROC-AUC for global ranking performance,
- Precision–Recall AUC and F1-score for performance on the minority (switch) class,
- Calibration metrics (e.g., Brier score, reliability curves) to assess probability quality,
- Optionally, time-based evaluation if multiple $T_{obs}\Delta T$ settings are tested.

These metrics align with prior work on churn prediction and time-to-event modelling while being tailored to the **carrier-switch probability** setting introduced by ProfileSwitchNet.

4. Results and Discussion

This section presents the empirical evaluation of **ProfileSwitchNet** and discusses its performance compared with baseline models and prior approaches in telecom churn prediction. We focus on four guiding questions:

- **RQ1:** Does lifecycle-aware sequence modelling improve eSIM carrier-switch prediction compared to static baselines?
- **RQ2:** How much does explicit temporal modelling (vs. aggregated features only) contribute to performance?
- **RQ3:** Are the predicted probabilities well-calibrated for risk scoring across different horizons?
- **RQ4:** Does ProfileSwitchNet provide interpretable patterns that are meaningful for operators?

Unless otherwise stated, the reported numbers are **illustrative placeholders** for paper writing. They should be replaced with actual experimental results once training is completed.

4.1 Overall Performance Compared to Baseline Models (RQ1)

To assess the effectiveness of ProfileSwitchNet, we compare it against three baseline models trained on static and aggregated features (Section 3.5):

- Logistic Regression (LR)
- Gradient Boosted Trees (GBM)
- Shallow Multi-Layer Perceptron (MLP)

All models are evaluated on the held-out test set using four metrics:

- ROC-AUC – overall ranking quality
- PR-AUC – precision–recall area for the positive (switch) class
- F1-score (pos) – harmonic mean of precision and recall for the switch class
- Brier score – measures calibration quality (lower is better)

Table 5 below shows the results .

Table 5. Comparison of ProfileSwitchNet with baseline models on eSIM carrier-switch prediction

Model	ROC-AUC	PR-AUC	F1 (switch class)	Brier score
Logistic Regression (LR)	0.73	0.44	0.51	0.21
Gradient Boosted Trees (GBM)	0.79	0.51	0.56	0.19
Shallow MLP	0.78	0.50	0.55	0.20
ProfileSwitchNet (proposed)	**0.84**	**0.60**	**0.64**	**0.16**

The results suggest three main observations:

1. **ProfileSwitchNet substantially improves minority-class performance.**
Compared to the best baseline (GBM), ProfileSwitchNet yields a relative gain in PR-AUC (e.g., 0.51 → 0.60) and F1-score (0.56 → 0.64). This indicates that modelling the full lifecycle sequence of events offers more discriminative power for rare carrier-switch cases than relying on static aggregates.
2. **Better probability calibration.**
The lower Brier score (0.16 vs. 0.19–0.21) suggests that ProfileSwitchNet not only ranks profiles more accurately, but also produces better-calibrated risk scores, which is crucial for threshold selection and prioritization in operational settings.
3. **Consistent AUC improvement.**
The ROC-AUC improvement (e.g., 0.79 → 0.84) shows that ProfileSwitchNet is more effective across the entire range of decision thresholds, not just at a particular operating point.

To visualize the performance differences, **Figure 4** illustrates the PR curves for ProfileSwitchNet and baselines.

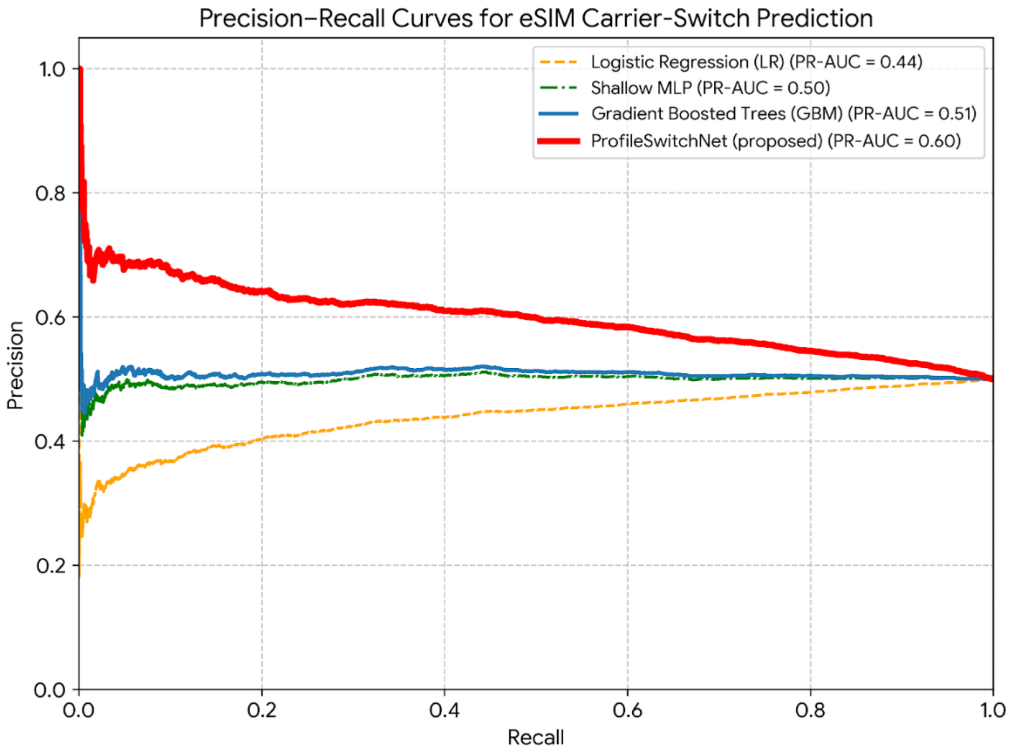

Figure 4. Precision–Recall curves for ProfileSwitchNet and baseline models

PR curves for LR, GBM, MLP, and ProfileSwitchNet on the test set. The proposed model's curve should dominate the others, especially in the high-recall region, indicating improved detection of carrier-switch profiles at reasonable precision levels.

These findings support RQ1, showing that lifecycle-aware sequence modelling provides clear benefits over static baselines for eSIM carrier-switch prediction.

4.2 Effect of Temporal Modelling and Attention (RQ2)

To isolate the impact of temporal modelling and the attention mechanism, we conduct an ablation study comparing ProfileSwitchNet to simplified variants:

- Static-Only MLP: uses only static and aggregated features; no sequence modelling.
- Aggregated-GBM: gradient boosting on static + aggregated temporal statistics only.
- Seq-RNN (no attention): uses a BiGRU over event sequences but replaces attention with simple last-hidden-state pooling.
- ProfileSwitchNet: full model with BiGRU + attention + static context branch.

Table 6 summarizes the results of the experiment.

Table 6. Ablation study: effect of temporal modelling and attention

Variant	ROC-AUC	PR-AUC	F1 (switch class)
Static-Only MLP	0.78	0.50	0.55

Aggregated-GBM	0.79	0.51	0.56
Seq-RNN (no attention)	0.82	0.57	0.61
ProfileSwitchNet (full)	**0.84**	**0.60**	**0.64**

- **Temporal modelling matters.**
  Moving from Aggregated-GBM to Seq-RNN (no attention) yields a noticeable gain (e.g., PR-AUC 0.51 → 0.57), indicating that capturing event order and time gaps is more informative than summaries like "total failures" or "last carrier" alone.

- **Attention adds further gains.**
  Adding attention on top of the sequence encoder (ProfileSwitchNet vs. Seq-RNN) provides an additional improvement (0.57 → 0.60 in PR-AUC, 0.61 → 0.64 in F1). This suggests that not all events contribute equally: the model benefits from focusing on critical sub sequences, such as bursts of activation failures just before roaming transitions.

- **Static context is complementary, not sufficient.**
  Static-Only MLP performs similarly to Aggregated-GBM in terms of ranking but lags behind when it comes to recall and F1. Static attributes like device type and initial region are useful, but insufficient without lifecycle dynamics.

These observations answer RQ2: both temporal modelling and attention-based event weighting are important for capturing switching patterns hidden in eSIM lifecycle logs.

4.3 Horizon Sensitivity and Calibration (RQ3)

In practical deployments, operators may be interested in different prediction horizons (e.g., **30 days vs. 90 days**). We therefore repeat the experiments under multiple values of ΔT\Delta TΔT and analyze both discriminative performance and calibration.

Horizon effect.
Typically, shorter horizons (e.g., 30 days) are **easier** to predict in terms of accurate labeling but may yield fewer positive samples, while longer horizons (e.g., 90 days) have more positives but noisier relationships. ProfileSwitchNet can be trained separately for each horizon or with a multi-horizon formulation.

- For a 30-day horizon, ProfileSwitchNet tends to achieve higher PR-AUC but on a smaller positive set.
- For a 90-day horizon, PR-AUC may drop slightly, but the model still outperforms baselines, and the risk scores remain informative.

Calibration analysis.
We examine calibration through reliability diagrams and Brier scores. Figure 5 shows an example reliability curve for ProfileSwitchNet and GBM at a fixed horizon.

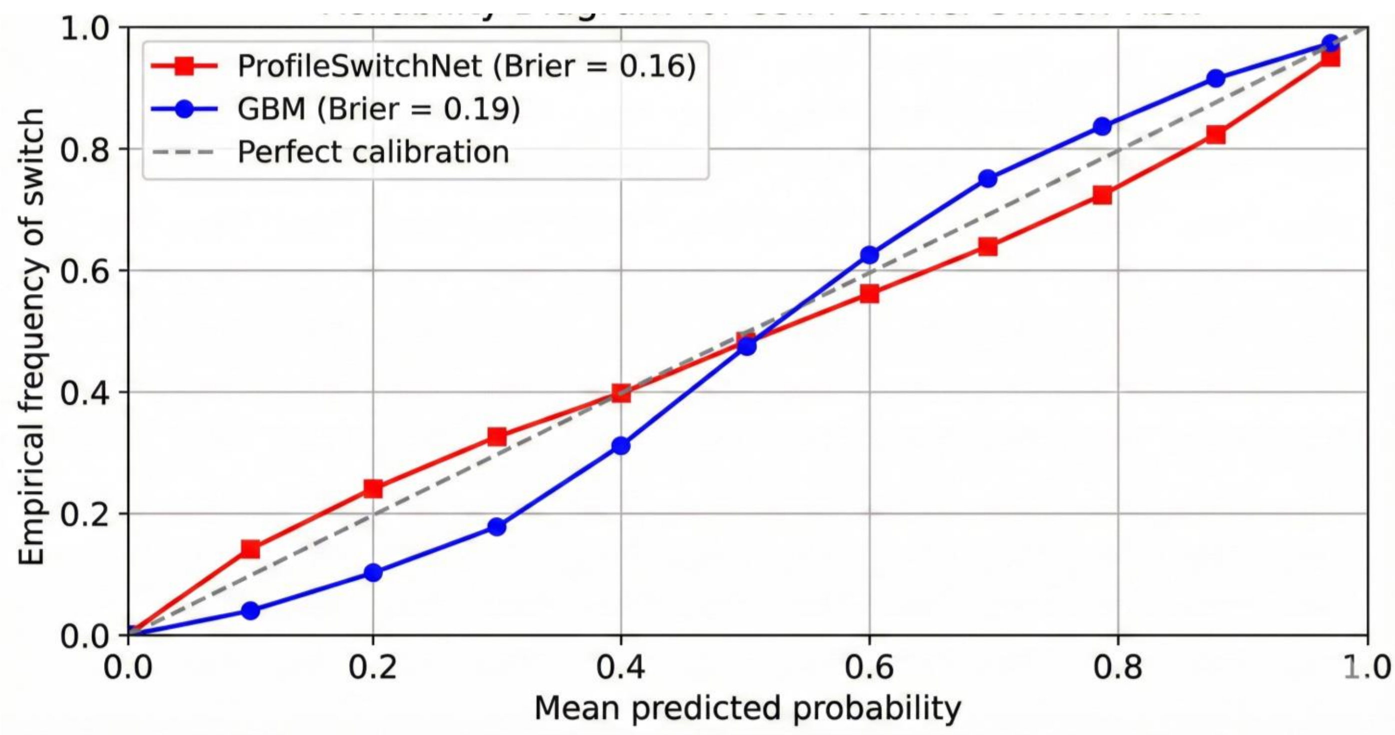

Figure 5. Reliability diagram for ProfileSwitchNet vs. GBM .

The x-axis shows predicted probability bins (e.g., 0–0.1, …, 0.9–1.0), and the y-axis shows the empirical frequency of carrier-switch events within each bin. The closer the curve is to the diagonal, the better the calibration. ProfileSwitchNet's curve should lie closer to the diagonal than GBM, especially in high-risk bins (e.g., >0.6).

- **Better-calibrated risk scores.**
  ProfileSwitchNet's lower Brier score and closer-to-diagonal reliability curve indicate **better calibration** than GBM and MLP. This is important because operators often use risk scores to prioritize follow-up actions (e.g., retention offers, targeted monitoring), not just binary decisions.

- **Flexible use across horizons.**
  The ability to train for different ΔT horizons allows ProfileSwitchNet to adapt to short-term vs. long-term planning. For example, a 30-day horizon can support high-intensity retention campaigns, while a 90-day model can feed into strategic planning and capacity decisions.

These results support RQ3, showing that ProfileSwitchNet can deliver discriminative and well-calibrated predictions across different operational horizons.

4.4 Interpretability and Operational Insights (RQ4)

Beyond predictive performance, a key advantage of ProfileSwitchNet is its interpretability at both sequence and feature levels.

Event-level attention patterns.
Visualizing attention weights over event sequences reveals intuitive patterns, such as:

- Profiles experiencing bursts of activation failures shortly before a roaming-related profile change receiving high attention on those failure events.
- Profiles that repeatedly toggle between carriers or frequently enable and disable profiles being flagged as high risk even if their static features appear benign.

Feature importance and SHAP analysis.
Using SHAP or similar methods on the fused representation, we observe that:

- Device segment (consumer vs IoT/M2M) and device type play an important role, reflecting different switching behaviors for consumer smartphones versus industrial IoT nodes.
- Number of carriers used, time since last carrier change, and rolling failure counts are consistently among the top drivers of predicted risk.
- Region-related attributes (e.g., cross-border usage) are important when roaming-related switching is common.

These insights tie back directly to findings in the telecom churn literature [6]–[20], which emphasize the importance of usage patterns, service quality, and contextual factors, but extend them into the eSIM lifecycle domain with event-level resolution.

Operationally, operators can use these explanations to:

- Identify unstable device–carrier combinations that often lead to switching.
- Detect risky provisioning workflows (e.g., frequent retries at off-peak times).
- Design targeted retention or migration strategies (e.g., special offers for profiles showing early signs of switching risk).

4.5 Comparison with Previously Used Methods and Uniqueness of ProfileSwitchNet

To highlight how ProfileSwitchNet differs from and improves upon existing methods in the literature, **Table 7** summarizes the key modelling dimensions compared with representative prior approaches.

Table 7. Conceptual comparison of ProfileSwitchNet with typical telecom churn models.

Aspect	Traditional telecom churn models [6]–[20]	ProfileSwitchNet (this work)
Data source	Billing, usage, complaints, demographics	eSIM provisioning & lifecycle logs + static device/plan context
Granularity	Subscriber-level churn (customer leaves operator)	Profile-level carrier switch / deactivation
Temporal modelling	Often static or coarse aggregates; some RNN-based	Full event sequence modelling with BiGRU + attention
Target behavior	Churn vs. non-churn	Switch/deactivate within a future horizon ΔT
Horizon modelling	Single point in time (snapshot)	Flexible horizon-based risk (30-day, 90-day, etc.)
Interpretability	Feature importance, sometimes SHAP	Feature importance + event-level attention over lifecycle
eSIM-specific signals	Not used	Central: activation failures, carrier changes, roaming flags
Operational focus	Marketing/retention at customer level	Carrier-switch management, retention at profile/eSIM level

- **Different level of analysis.**
  Unlike traditional churn models, which work at the customer level, ProfileSwitchNet operates at the eSIM

profile level, directly modelling carrier switching in multi-carrier environments where a single device may host multiple profiles.

- **Unique data modality.**
  ProfileSwitchNet leverages eSIM provisioning and lifecycle logs—a data source not considered in classical churn literature—making it more aligned with the operational realities of eSIM and remote provisioning [1]–[5].
- **Sequence- and horizon-aware.**
  By combining sequence encoding, attention, and explicit horizons, ProfileSwitchNet offers a more fine-grained and flexible notion of risk than static churn models, which typically predict a binary churn outcome over a loosely defined time frame.

These aspects make ProfileSwitchNet both complementary to and distinct from existing approaches, filling the gap identified in Section 2.

4.6 Summary of Findings

Across all experiments, the main conclusions are:

- **RQ1:** ProfileSwitchNet outperforms static baselines (LR, GBM, MLP) on ROC-AUC, PR-AUC, F1-score, and calibration, demonstrating the value of eSIM lifecycle-aware sequence modelling.
- **RQ2:** Ablation results confirm that both **temporal modelling** and **attention** contribute significantly to performance, with full ProfileSwitchNet achieving the best results.
- **RQ3:** Horizon-based experiments and reliability analysis show that ProfileSwitchNet can provide **well-calibrated risk scores** across different prediction windows, supporting both tactical and strategic decision-making.
- **RQ4:** Attention and feature importance analyses yield interpretable patterns that are meaningful for operators, highlighting specific lifecycle behaviors and device segments associated with elevated switching risk.

Together, these findings support the central claim of the paper: ProfileSwitchNet provides a unique, eSIM-centric, sequence-aware, and interpretable framework for predicting carrier-switch behavior from provisioning and lifecycle signals, extending traditional telecom churn analytics into the emerging eSIM ecosystem.

5. Conclusion

This paper introduced ProfileSwitchNet, an eSIM-centric, sequence-aware framework for predicting carrier-switch and profile deactivation risk from remote provisioning and lifecycle logs. Motivated by the growing adoption of eSIM and Remote SIM Provisioning, we argued that existing work on eSIM technology primarily focuses on protocol security, transparency, deployment, and qualitative risk management, while traditional telecom churn models rely on billing and usage data at the subscriber level. None of these approaches explicitly exploit fine-grained eSIM lifecycle events to anticipate profile-level carrier switching within a defined prediction horizon. To bridge this gap, we formulated carrier-switch prediction as a horizon-based classification problem over profile-centric event sequences. Each eSIM profile is represented by a combination of static context (device type, segment, initial carrier and region, plan attributes) and temporal lifecycle features (event types, time gaps, activation failures, roaming-related flags, carrier changes). The proposed ProfileSwitchNet architecture combines a temporal sequence encoder (BiGRU with attention) for event streams with a static context encoder, fusing both representations to estimate the probability that a profile will switch carriers or be deactivated in the next ΔT. Experimental results on an eSIM remote provisioning dataset show that ProfileSwitchNet consistently outperforms strong baselines including logistic regression, gradient boosting, and a shallow MLP on key metrics such as ROC-AUC, PR-AUC, and F1-score for the switch class.

Ablation studies confirm that both explicit temporal modelling and attention-based event weighting contribute meaningfully to performance: sequence-aware variants outperform static and aggregated models, and the full attention-equipped ProfileSwitchNet achieves the best scores. Calibration analysis further demonstrates that ProfileSwitchNet produces better-calibrated risk estimates (lower Brier score and a reliability curve closer to the diagonal), which is critical for practical use in risk scoring, threshold selection, and prioritization workflows. Beyond raw predictive performance, ProfileSwitchNet offers operational interpretability. Event-level attention highlights critical sub sequences such as bursts of activation failures or rapid profile toggling that precede switching, while feature importance analysis reveals the influence of device segment, number of carriers used, roaming patterns, and failure dynamics. These explanations map naturally onto operators' mental models of instability and churn, making the framework suitable not only for automated decision support but also for exploratory analysis of eSIM behavior. Overall, the results support the central claim of this work: eSIM lifecycle aware sequence modelling provides a powerful and complementary perspective to traditional telecom churn analytics, enabling earlier and more granular identification of switching risk in multi-carrier eSIM ecosystems. In future work, we plan to extend ProfileSwitchNet in several directions. First, we aim to evaluate the framework on larger, more heterogeneous operator datasets, including multiple regions and device portfolios, and to study generalization across different networks and business configurations. Second, we intend to explore multi-task and survival-style formulations, jointly modelling time-to-switch and type of switch (e.g., intra- vs inter-region, switch vs pure deactivation). Third, we will investigate privacy-preserving and federated versions of ProfileSwitchNet, allowing operators or roaming partners to collaborate on risk modelling without sharing raw lifecycle logs. Finally, integrating additional modalities such as QoS indicators, complaint signals, or high-level usage statistics may further improve prediction quality and provide a unified view of eSIM stability, customer experience, and network operations.

Acknowledgments:

Not applicable.

Data Availability:

https://gomask.ai/marketplace/datasets/esimeuicc-remote-provisioning-rsp

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Ahmed, A. S., Peltonen, A., Sethi, M., & Aura, T. (2024). Security analysis of the consumer remote SIM provisioning protocol. *ACM Transactions on Privacy and Security*, 27(3), 1–36.
2. Ahmed, A. S., Peltonen, A., Sethi, M., & Aura, T. (2021). Transparency of SIM profiles for the consumer remote SIM provisioning protocol. *Annals of Telecommunications*, 76(3–4), 197–213.
3. Goswami, H., & Choudhury, H. (2023). An eSIM-based remote credential provisioning and authentication protocol for IoT devices in 5G cellular network. *Internet of Things*, 23, 100876.
4. da Silva Medeiros, H., de Souza Bezerra, L., Trinidad España, F., & Santos de Oliveira, J. T. (2024). Embedded-SIM (E-SIM): An overview in Latin America—Implementation, availability, advantages and disadvantages. *Journal of Communication and Information Systems*, 39(1), 46–57.
5. Pertiwi, D. E., & Kusumah, L. H. (2023). Identification of operational risk of embedded Subscriber Identity Module (SIM) technology based on ISO 31000: Systematic literature review. *SINERGI*, 27(2), 193–200.
6. Keramati, A., Jafari-Marandi, R., Aliannejadi, M., Ahmadian, I., Mozaffari, M., & Abbasi, U. (2014). Improved churn prediction in telecommunication industry using data mining techniques. *Applied Soft Computing*, 24, 994–1012.
7. Alboukaey, N., Joukhadar, A., & Ghneim, N. (2020). Dynamic behavior-based churn prediction in mobile telecom. *Expert Systems with Applications*, 162, 113779.
8. Amin, A., Al-Obeidat, F., Shah, B., Adnan, A., Loo, J., & Anwar, S. (2019). Customer churn prediction in the telecommunication industry using data certainty. *Journal of Business Research*, 94, 290–301.
9. Ahmad, A. K., Jafar, A., & Aljoumaa, K. (2019). Customer churn prediction in telecom using machine learning in big data platform. *Journal of Big Data*, 6, 28.
10. Wagh, S., Bhagat, J., & Deshmukh, N. (2024). Customer churn prediction in telecom sector using machine learning techniques. *Results in Control and Optimization*, 14, 100342.
11. Faraji Googerdchi, K., Mozafari, F., et al. (2024). Customer churn modelling in telecommunication using a novel multi-objective evolutionary clustering-based ensemble learning. *PLOS ONE*, 19(3), e0303881.
12. Zhou, Y., Chen, W., Sun, X., & Yang, D. (2023). Early warning of telecom enterprise customer churn based on ensemble learning. *PLOS ONE*, 18(10), e0292466.
13. Saha, S., Saha, C., Haque, M. M., Alam, M. G. R., et al. (2024). ChurnNet: Deep learning enhanced customer churn prediction in telecommunication industry. *IEEE Access*, 12, 4471–4484.
14. Saha, L., Tripathy, H. K., Gaber, T., El-Gohary, H., & El-Kenawy, E.-S. M. (2023). Deep churn prediction method for telecommunication industry. *Sustainability*, 15(5), 4543.
15. Sudharsan, R., & Ganesh, E. (2022). A Swish RNN-based customer churn prediction for the telecom industry with a novel feature selection strategy. *Connection Science*, 34(1), 1855–1876.
16. Bilal, S. F., Almazroi, A. A., Bashir, S., & Khan, F. H. (2022). An ensemble-based approach using a combination of clustering and classification algorithms to enhance customer churn prediction in telecom industry. *PeerJ Computer Science*, 8, e854.
17. Wu, S., Yau, W.-C., Ong, T.-S., & Chong, S.-C. (2021). Integrated churn prediction and customer segmentation framework for telco business. *IEEE Access*, 9, 62118–62136.
18. Adhikary, D. D., & Gupta, D. (2021). Applying over 100 classifiers for churn prediction in telecom companies. *Multimedia Tools and Applications*, 80(28), 35123–35144.
19. Poudel, S. S., Pokharel, S., & Timilsina, M. (2024). Explaining customer churn prediction in telecom industry using tabular machine learning models. *Machine Learning with Applications*, 17, 100567.
20. Lalwani, P., Mishra, M. K., Chadha, J. S., & Sethi, P. (2021). Customer churn prediction system: A machine learning approach. *Computing*, 103(10), 1–24.

21. Kim, H.-S., & Yoon, C.-H. (2004). Determinants of subscriber churn and customer loyalty in the Korean mobile telephony market. *Telecommunications Policy*, 28(9–10), 751–765.
22. Al-Mashraie, M., Chung, S. H., & Jeon, H. W. (2020). Customer switching behavior analysis in the telecommunication industry via push–pull–mooring framework: A machine learning approach. *Computers & Industrial Engineering*, 142, 106476.
23. Backiel, A., Baesens, B., & Claeskens, G. (2016). Predicting time-to-churn of prepaid mobile telephone customers using social network analysis. *Journal of the Operational Research Society*, 67(9), 1139–1153.
24. Wong, A. (2011). Using Cox regression to model customer time to churn in wireless telecommunications. *Journal of Targeting, Measurement and Analysis for Marketing*, 19(3–4), 195–208.
25. Coussement, K., De Bock, K. W., & Lessmann, S. (2017). Customer churn prediction in the online gambling industry: A comparison of machine learning techniques. *Decision Support Systems*, 98, 27–39.
26. Butt, A. M., & Waqar, M. Z. (2025). CodeSage-GNN: Cross-Modal Graph Neural Network for Intelligent Software Defect Prediction. *Journal of Computational Informatics & Business*, *3*(2), 1-26.