CareNet: Linking Home-router Network Traffic to DSM-5 Depressive Behavior Indicators

Stephan Nef, Bruno Rodrigues
Embedded Sensing Group ESG, School of Computer Science SCS, University of St. Gallen HSG, Switzerland
E-mail: stephan.nef@student.unisg.ch|bruno.rodrigues@unisg.ch

***Abstract*—Digital mental-health sensing increasingly depends on mobile or wearable devices that require intrusive permissions and continuous user compliance. We present *CareNet*, a router-centric system that transforms household network metadata into interpretable behavioral indicators aligned with DSM-5 depressive-symptom domains. All processing occurs locally at the home gateway, preserving privacy while maintaining visibility of temporal routines.**

The core contribution is the *Fuzzy Additive Symptom Likelihood (FASL)*, a transparent formulation that fuses header-level metrics into daily criterion-level likelihoods using bounded fuzzy memberships and additive aggregation. Combined with a DSM-style temporal gate, FASL integrates short-term traffic fluctuations into persistent, clinically interpretable indicators. Evaluation on realistic multi-day traces shows that CareNet captures characteristic patterns such as delayed sleep timing and attentional instability without payload inspection. The results highlight the feasibility of reproducible, explainable behavioral inference from router-side telemetry.

***Index Terms*—Router-centric sensing, network metadata, DSM-5, Behavior Observation Metrics**

I. Introduction

Depression is a leading contributor to the global burden of disease and often begins in adolescence, where it predicts lower educational attainment and poorer adult outcomes [15], [5]. Early intervention matters, yet delays from symptom onset to evidence-based care are common and worsen prognosis [19], [15]. Digital behavioral monitoring has linked passively collected signals to sleep, activity, and social rhythm, but device-centric approaches require apps, ongoing compliance, and invasive permissions; in multi-person households they are unreliable and raise privacy and deployment concerns [22], [27], [29].

This paper investigates to which extent a router-centric vantage point can provide meaningful insights on mental-health awareness. The central idea (*cf.* in Figure 1) is to analyze header-level timing and coarse destination information (*e.g.*, Public Suffix List and IAB - Interactive Advertising Bureau - content taxonomy labels) without inspecting payloads, and to summarize these signals into *Behavior Observation Metrics (BOMs)* that align with symptom patterns described in DSM-5 (Diagnostic and Statistical Manual of Mental Disorders) *e.g.*, disturbed sleep–wake rhythm, social withdrawal [16], [10], [2]. All processing remains local to avoid data exposure and operational overhead.

We consider a multi-person household in which all devices share a single home gateway. In this setting, we propose *CareNet*, an end-to-end pipeline (*cf.* Figure 1) that: (i) captures minute level network activity and enriches flows with coarse domain labels; (ii) derives interpretable features such as first online time, nocturnal bursts, protocol diversity, and repetition from packet headers only; (iii) groups related features into *Behavior Observation Metrics* (BOMs) with explicit *risk direction* and clinical rationale; (iv) fuses these into *criterion level likelihoods* through transparent fuzzy memberships; and (v) summarizes rolling histories using a DSM-style gate to produce day level indicators. The design favors transparency and auditability over black-box accuracy and confines observation to non content metadata for privacy preserving and interpretable insight.

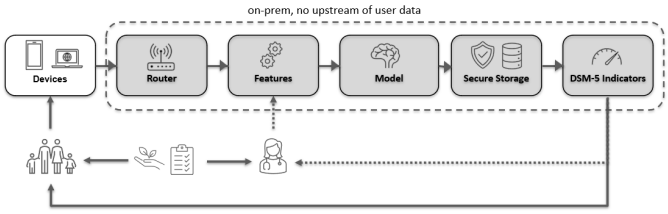

Fig. 1: Overview of CareNet's approach: From parsing network traffic and extracting features to mapping features into BOMs and DSM-aligned indicators. CareNet's source-code is available: [18]

Four main challenges arise in this scenario. The first is *temporal dilution*, where short but behaviorally significant deviations (*e.g.*, an unusually late night of activity or an evening of social withdrawal) vanish when data are averaged into daily or weekly aggregates [22], [28], [8]. The second is *multi-user mixtures*, since household traffic merges multiple devices and users, obscuring individual changes [30], [14]. The third concern *underused header dynamics*: router-level studies often reduce traffic to byte counts, overlooking informative timing, burst, and destination patterns that could reveal behavior without content inspection [30]. Finally, the fourth challenge is the *adoption barrier* of opaque models, as lack of interpretability hinders clinical trust and calls for transparent mappings from features to DSM-aligned indicators [26], [20].

CareNet makes the following contributions:

- **Router-centric telemetry:** a household-gateway vantage that extracts behavioral signals from header metadata while preserving privacy.
- **Behavior Observation Metrics (BOMs):** a criterion-

Related	Short description	Router	Meta	Attr.	Interp.	DSM gate	Privacy
Wang *et al.* (2014) [27]	Smartphone app (sleep, mobility, social rhythm)	○	○	◐	◐	○	○
Saeb *et al.* (2015) [22]	Phone GPS/accelerometer for PHQ-9 correlation	○	○	◐	○	○	○
Wang *et al.* (2018) [28]	Phone + Fitbit multimodal fusion	○	◐	◐	◐	○	◐
Saeb *et al.* (2017) [21]	Mobile passive sleep monitoring	○	○	◐	○	○	○
Kushlev *et al.* (2016) [13]	Notification bursts and attention effects	○	◐	◐	◐	○	◐
Alamoudi *et al.* (2024) [1]	Sleep-tracker app, anxiety/depression prediction	○	○	◐	◐	○	○
Cacheda *et al.* (2019) [4]	Social-media text mining (Twitter)	○	○	◐	○	○	○
De Choudhury *et al.* (2016) [6]	Reddit linguistic shift analysis	○	○	◐	◐	○	○
Bucur *et al.* (2024) [3]	LLM-based social-media depression detection	○	○	◐	○	○	○
Ware *et al.* (2018) [30]	Campus Wi-Fi metadata for screening	◐	●	◐	◐	○	◐
Mammen *et al.* (2021) [14]	Home-Wi-Fi sleep duration inference	●	●	◐	◐	○	●
Ware *et al.* (2020) [29]	Smartphone network usage predicting mood	●	●	○	◐	○	●
Stachl *et al.* (2020) [24]	Large-scale smartphone metadata, personality traits	○	◐	◐	◐	○	◐
CareNet (this paper)	Router-centric metadata for behavioral indicators	●	●	●	●	●	●

TABLE I: Representative related work and their alignment with categories that are central. Legend: ● = provided; ◐ = partially; ○ = not provided.

aligned mapping from network timing and domain diversity to interpretable behavioral indicators.

- **Fuzzy Additive Symptom Likelihood (FASL):** a transparent fusion model combining fuzzy memberships and a DSM-style gate to derive daily DSM-5–aligned likelihoods.
- **Prototype and evaluation:** an end-to-end implementation on real traces demonstrating feasibility, stability, and reproducibility under privacy constraints.

Our target setting is adolescents in multi-person households. Indicators are *non-diagnostic* and intended to support early observation by caregivers and clinicians; payloads are not inspected, and processing can remain local. The work complements (not replaces) established tools such as PHQ-9 [12], [23].

The paper is organized as follows. Section II reviews background and related work. Section IV details the *CareNet* design and FASL formulation. Section V reports the evaluation, and Section VII concludes the paper.

II. Background: DSM-5 and BOM

DSM-5 defines a *Major Depressive Disorder (MDD)* as at least five of nine core symptoms persisting for a minimum of two weeks, with one being either depressed mood or loss of interest [26], [2]. *CareNet* follows this definition to align all behavioral indicators with clinically recognized criteria while preserving privacy and interpretability (Table II).

TABLE II: DSM-5 symptom indicators for MDD [2].

ID	Symptom Indicator
(1)	Depressed mood
(2)	Loss of interest or pleasure (anhedonia)
(3)	Significant weight change or appetite disturbance
(4)	Insomnia or hypersomnia
(5)	Psychomotor retardation or agitation
(6)	Fatigue or loss of energy
(7)	Feelings of worthlessness or excessive guilt
(8)	Diminished ability to think, concentrate, or indecisiveness
(9)	Recurrent thoughts of death or suicidality

Behavior Observation Metric (BOM) acts as a digital behavioral biomarker that aggregates observable network and environmental features into interpretable indicators aligned with DSM-5 symptom domains [22], [27]. While not diagnostic, it translates feature-level evidence into criterion-level likelihoods for early awareness. Each BOM emphasizes *validity* (explicit mapping to DSM-5 constructs), *reliability* (temporal stability and sensitivity to change), and *transparency* (auditable memberships, weights, and gate parameters within the framework).

III. Related work

Table I provides a summary of each related work, highlighting the focus or modality. Device-centric studies using apps or wearables show digital traces correlate with sleep, mobility, and mood [27], [22], [28], [21]. However, they score low on router-centric and metadata-only axes due to reliance on per-device instrumentation and permissions. Person attribution assumes one-device–one-user, complicating issues like short behavioral excursions being diluted by long windows and data fragility with low compliance. Models are either partly interpretable or opaque, and lack a formal DSM-style aggregation step.

A second line of work analyzes social-media content to detect depressive language [4], [6], [3]. These approaches achieve strong performance in text domains, but by design they violate the metadata-only constraint and carry privacy trade-offs; interpretability varies and criterion-level decision logic is typically absent. As a result, they are not suitable for a home-gateway that must avoid payload inspection.

Network-centric environmental sensing aligns with our goals, as WiFi studies [30], [17], [14], [29] achieve a full circle in metadata-only sensing and router vantage progress, revealing sleep and mobility patterns via network signatures. These studies, however, typically focus on aggregate households or single behaviors, leaving person-level attribution unresolved, interpretability partial, and lacking an explicit DSM-style gate. Additionally, the underutilization of rich

header dynamics like burst structure, protocol mix, and repetition limits the clinical feature mapping.

CareNet tackles challenges in various categories such as temporal dilution and multi-user mixtures by using a router-centric, metadata-only perspective. It introduces person-level attribution at gateways, enhances transparency through interpretable *Behavior Observation Metrics*, and uses a DSM-style gate for daily likelihood summaries while maintaining privacy. This approach complements existing device-centric and content-based methods.

IV. Methodology

CareNet is designed as a router-centric, privacy-first system that interprets network metadata within the household to infer behavioral metrics relevant to mental-health awareness. Figure 2 summarizes our approach and the closed information-flow loop that makes router-centric observation useful over time.

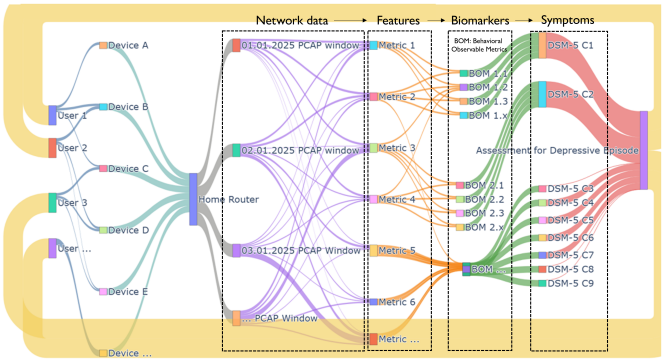

Fig. 2: Sankey diagram illustrating CareNet's general workflow

Figure 3 illustrates the end-to-end workflow. Traffic from all devices passes through the *Home Router* to the *Gateway*, where *Network Capture* records header metadata, the *ETL Parser* (Extract, Transform, Load) extracts key fields, and *Partitioned Storage* organizes them into short, fixed segments for local, privacy-preserving analysis.

In parallel (*cf.* Figure 3 right upper side), the system maintains an inventory of devices via *Device & Session Discovery* and resolves household ambiguity through *Identity Resolution (IP to Profile allocation)*, storing assignments in *User Profile Registry*. These identity services feed the analytical path at two points. First, they inform *Feature Extraction* so that the measurements per-segment reflect the intended individual, and second, condition *Feature Store*, where the records per-segment are retained. A lightweight *Feature Cache* accelerates iterative analysis and a *Parameter Registry* maintains analyst-controlled settings for memberships, weights, and gate parameters.

From the feature store, the measures aligned with the criterion are calculated in *Feature Engineering* and grouped into interpretable components within `BOM`. The *DSM-5 MDD Criterion Modules (C1-C9)* transform these components into daily probables based on a possible criterion

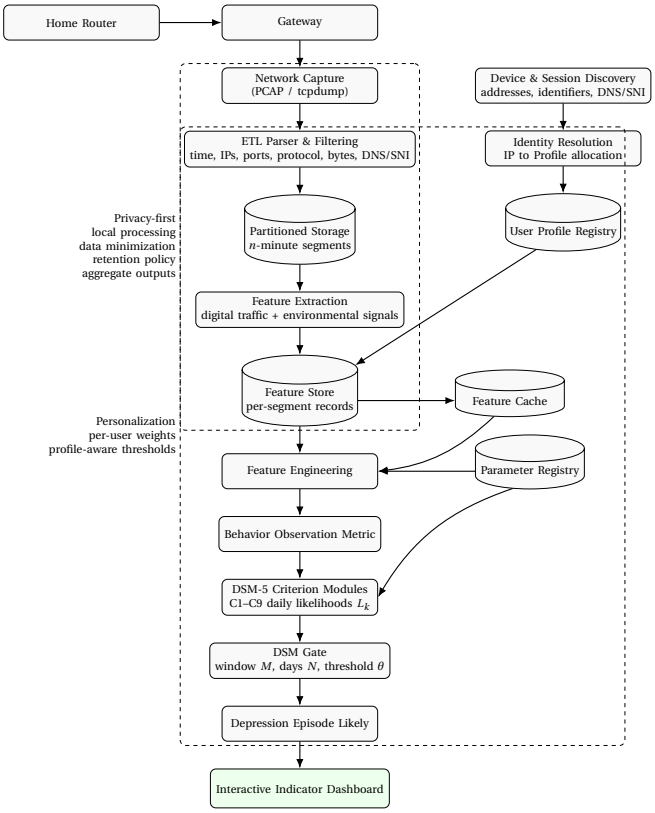

Fig. 3: *CareNet* detailed workflow, detailing key components and the flow through parsing, storage, feature extraction, and BOM aggregation toward DSM-5 MDD indicators and an interpretable gate.

$L_k \in [0,1]$. The *DSM Gate* then applies a clinically motivated rule on a rolling window M: a criterion is marked present when $L_k \geq \theta$ on at least N days.

A. Mapping Network Features to Behavioral Metrics

Table III summarizes how (selected) router-visible network features are grouped into interpretable *Behavior Observation Metrics (BOM)* and aligned with DSM-5 symptom domains. Each DSM-5 criterion (C1–C9) (*cf.* Table II on DSM-5 MDD symptoms) is bolstered by a concise array of criteria-aligned metrics that convert router-side measurements into interpretable signaling indicators with a clear orientation (*pro* or *contra*).

System-design-wise, metrics are intermediate objects in a structured mapping: related metrics are grouped into clinically meaningful BOM, and BOM are mapped to criterion-level likelihoods. This separation (metrics → BOM → criterion) keeps clinical meaning visible, enables profile-aware tuning through a parameter registry, and makes the pathway from quantitative network data to qualitative constructs auditable.

Network features are computed inside the privacy boundary from header/timing metadata only; registrable domains use eTLD+1 (Effective Top-Level Domain plus one), derived from the Public Suffix List (PSL) [16], [10], and optional top-

TABLE III: Feature mapping specification with descriptor/type.

Network feature (selected examples)	Biomarker (BOM)	Ind.	Rel.	Dir.	W	Desc.	Signal group
Distinct social domains (eTLD+1)	Reduced social interaction	(1)	Gradual	–	1.0	HLD	Service
Median chat session duration	Reduced social interaction	(1)	Gradual	–	1.0	LLD	Session
Long streaming session duration	Passive media binge	(1)	Gradual	+	1.0	LLD	Session
Down/Up byte ratio (D/U)	Passive media binge	(1)	Gradual	+	1.0	HLD	Volume
Revisit diversity (Shannon index)	Rumination loops	(1)	Gradual	–	1.0	HLD	Navigation
Short-interval repeated visits	Rumination loops	(1)	Gradual	+	1.0	LLD	Navigation
Hourly traffic coefficient of variation	Flattened diurnal rhythm	(1)	Gradual	–	1.0	HLD	Rhythm
Night/Day traffic ratio	Flattened diurnal rhythm	(1)	Gradual	+	1.0	HLD	Rhythm
Domain diversity (count eTLD+1)	Interest breadth	(2)	Gradual	–	1.0	HLD	Service
Category entropy (IAB topic mix)	Interest breadth	(2)	Gradual	–	1.0	HLD	Service
Chat session count (per day)	Engagement	(2)	Gradual	–	1.0	HLD	Session
Reply latency (chat / messaging)	Engagement	(2)	Gradual	+	1.0	LLD	Interaction
Passive/active ratio (D/U vs. upload bursts)	Passivity	(2)	Gradual	+	1.0	HLD	Volume
Upload-burst rate per active hour	Passivity	(2)	Gradual	–	1.0	HLD	Interaction
Delivery activity deviation (z-score)	Ordering activity	(3)	Gradual	both	1.0	HLD	Service
Late-night delivery share (>23:00)	Ordering timing	(3)	Gradual	+	1.0	HLD	Rhythm
Diet/nutrition domain exposure	Diet focus	(3)	Gradual	both	1.0	HLD	Service
Digital sleep onset (min after 22:00)	Shifted sleep timing	(4)	Gradual	+	1.0	HLD	Rhythm
Onset-time variability (14-day circular SD)	Shifted timing irregularity	(4)	Gradual	+	1.0	HLD	Rhythm
Main nightly idle-gap length	Sleep duration change	(4)	Gradual	both	1.0	HLD	Rhythm
Nocturnal micro-session count	Sleep fragmentation	(4)	Gradual	+	1.0	LLD	Rhythm
Inter-awakening gap (median)	Sleep fragmentation	(4)	Gradual	–	1.0	LLD	Rhythm
Daytime idle ratio (08–18 h)	Daytime hypersomnia / flattening	(4)	Gradual	+	1.0	HLD	Rhythm
Night/Day traffic ratio	Flattened rhythm	(4)	Gradual	+	1.0	HLD	Rhythm
Wi-Fi re-associations/DHCP renewals (count)	Restlessness (device churn)	(5)	Gradual	+	1.0	LLD	Mgmt
Very short sessions (< 15 s) count	Restlessness (micro-activity)	(5)	Gradual	+	1.0	LLD	Session
Median inter-session gap	Motor slowing vs. agitation	(5)	Gradual	both	1.0	HLD	Session
Midday idle minutes share	Low energy	(6)	Gradual	+	1.0	HLD	Rhythm
Daytime session count	Low energy	(6)	Gradual	–	1.0	HLD	Session
Sent bytes per active hour	Effortful interaction	(6)	Gradual	–	1.0	HLD	Volume
Upload-burst rate per active hour	Effortful interaction	(6)	Gradual	–	1.0	HLD	Interaction
Inter-click interval (mean/variance)	Slow browsing tempo	(6)	Gradual	+	1.0	LLD	Interaction
Mental-health resource domains (visits)	Help-seeking / self-worth	(7)	Gradual	+	1.0	HLD	Service
Therapist directory domains (visits)	Help-seeking	(7)	Gradual	+	1.0	HLD	Service
Median page dwell time	Fragmented focus	(8)	Gradual	–	1.0	LLD	Navigation
DNS lookup burst rate (tab-hopping)	Fragmented focus	(8)	Gradual	+	1.0	LLD	Navigation
Notification-triggered micro-sessions	Interruptions	(8)	Gradual	+	1.0	LLD	Session
Back-navigation rate after click	Indecisive search	(8)	Gradual	+	1.0	HLD	Navigation
Search back-to-landing within 60 s	Indecisive search	(8)	Gradual	+	1.0	HLD	Navigation
SERP time-to-first-click (domain change)	Indecision latency	(8)	Gradual	+	1.0	LLD	Navigation
Crisis-line domains (visits)	Crisis seeking	(9)	Gradual	+	1.0	HLD	Service
Self-harm community domains (visits)	Self-harm exposure	(9)	Gradual	+	1.0	HLD	Service
Cloud-backup surge (GB/day)	Digital affairs / planning	(9)	Gradual	+	1.0	HLD	Volume

Notes: Ind. numbers = DSM-5 MDD criteria (1–9). Dir. ∈ {+, −, both} denotes directionality (positive, negative, both).
HLD = daily/high-level aggregate; LLD = low-level/session or burst measure. Service grouping can use eTLD+1 via PSL and optional IAB topic labels (header-only; content not inspected).

ical labels are header-only (*cf.* Table III notes). All metrics are daily aggregates or summaries produced from n-minute partitions and written to the Feature Store, enabling local, profile-aware processing via the Parameter Registry.

Each feature x_i is first normalized against the user's personal baseline to obtain a deviation $z_{i,t}$, and then passed through a monotonic mapping function $g_i(\cdot)$ that expresses the directionality of the feature (*pro* or *contra*) with respect to the behavioral construct.

B. Fuzzy Additive Symptom Likelihood (FASL)

To convert packet-level network features into interpretable DSM-5 MDD indicators, we introduce the *Fuzzy Additive Symptom Likelihood* (FASL). FASL treats traffic-derived metrics (*e.g.*, flow duration, idle-gap length, byte ratios) as noisy network measurements and fuses them into daily, criterion-level likelihoods. All this process is done at the edge (gateway) and realized in the steps as follows.

(1) Triangular memberships. Each network metric $x_i(t)$ obtained from router telemetry or session statistics is mapped to $[0,1]$ using a simple triangular membership $\mu_i(\cdot)$:

$$\mu_{\text{tri}}(x;\ell,m,h)=\begin{cases}0, & x\le \ell,\\ \dfrac{x-\ell}{m-\ell}, & \ell<x\le m,\\ \dfrac{h-x}{h-m}, & m<x<h,\\ 0, & x\ge h.\end{cases} \tag{1}$$

The triangular membership is an ideal fit as it offers a computationally efficient and fully interpretable mapping with finite support in $[\ell,h]$, which is essential for continuous edge-level processing on the gateway. The function transforms a normalized network metric (*e.g.*, packet-rate deviation) into bounded evidence with finite support, making threshold logic explicit and verifiable.

(2) Biomarkers or Behavioral Observable Metrics. Related network features form a signed component whose evidence remains bounded and interpretable:

$$S_{k,b}(t)=\text{clip}_{[-1,1]}\left(\sum_{i\in\mathscr{I}_b} w_{b,i}\, s_{k,i}\, \mu_{k,i}\big(x_i(t)\big)\right),\quad \sum_i w_{b,i}=1. \tag{2}$$

Features assigned to behavioral component b (for DSM-5 criterion k) are fused additively with normalized weights. Packet-level metrics from the same traffic dimension (*e.g.*, session timing or domain diversity) are fused using normalized weights, where $s_{k,i}\in\{+1,-1\}$ encodes pro/contra orientation so that upload scarcity and burstiness, for example, can have opposite signs.

(3) Criterion likelihood. Signed component evidence $S_{k,b}\in[-1,1]$ is first mapped to a non-negative domain and then fused into a daily likelihood:

$$\hat{S}_{k,b}(t)=\tfrac{1}{2}\big(S_{k,b}(t)+1\big)\in[0,1], \tag{3}$$

$$L_k(t)=\text{clip}_{[0,1]}\left(\sum_{b\in\mathscr{B}_k}\nu_{k,b}\,\hat{S}_{k,b}(t)\right),\qquad \sum_b \nu_{k,b}=1. \tag{4}$$

The affine map $\frac{1}{2}(S+1)$ prevents silent negative-sum collapses and preserves interpretability as evidence mass.

(4) DSM-style gate. A day is positive for criterion k if the likelihood exceeds a threshold; presence requires at least N positives in the last M days:

$$I_k(d)=\mathbb{I}[\,L_k(d)\ge\theta\,],\qquad \text{present}_k(t)=\mathbb{I}\left[\sum_{d=t-M+1}^{t} I_k(d)\ge N\right]. \tag{5}$$

The gate applies a temporal sliding window over daily likelihoods, which is analogous to a moving-average filter in network monitoring—to require that signals persist over N of M days before being flagged.

(5) Episode rule. Following DSM-5, an episode is likely when ≥ 5 criteria are present, including ≥ 1 core (C1 or C2):

$$\text{epi}(t)=\mathbb{I}\left[\sum_{k=1}^{9}\text{present}_k(t)\ge 5\ \wedge\ \exists c\in\{1,2\}:\text{present}_c(t)=1\right]. \tag{6}$$

This gating logic operationalizes the DSM-5 "nearly every day" rule through a moving window over daily network-derived likelihoods, analogous to sliding-window detection in traffic analytics. When at least five DSM-5 criteria, including one core criterion, are active, the router reports an *episode-likely* state. All parameters (w, θ, M, N) are explicit and tunable, providing the same transparency and reproducibility expected in network-telemetry pipelines.

The resulting decision process can be visualized as a sliding window over daily likelihoods, where each day is marked positive or negative depending on whether its criterion likelihood exceeds the threshold θ. Figure 4 illustrates this temporal gating logic and how it maps continuous router-derived signals to discrete DSM-5 criterion presence.

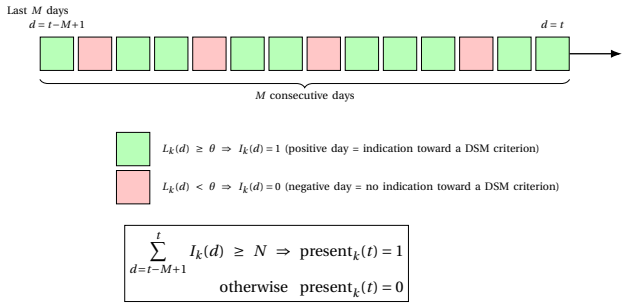

Fig. 4: Each square represents *one day* in the last M days. A day is green if the daily likelihood for criterion k exceeds the threshold ($L_k(d) \geq \theta$), otherwise red. Count the green (positive) days and if at least N of the last M days are positive, the criterion is marked present.

Days with insufficient evidence are excluded from the window sum (data-quality rule), ensuring that decisions are not driven by missing or unreliable measurements. In practice, this design allows a household gateway to behave like an interpretable measurement node, aggregating packet header-level observations, applying fuzzy inference locally, and exposing clear, verifiable indicators rather than opaque scores (such as using an LLM to link features to BOMs, and BOMs to symptoms).

V. Experimental Evaluation

This section evaluates whether the router-centric prototype can (i) extract stable traffic header-level features from raw PCAP, (ii) transform them into day-level metrics via fixed windowing and transparent membership functions, and (iii) aggregate them into DSM-aligned indicators using a persistence gate.

A. Setup and datasets

Hardware. For implementation and testing, *CareNet* was evaluated on a Lenovo x64 (AMD Ryzen 7 7730U), containerized with Docker [7] for portability, and deployed to Streamlit Community Cloud (https://unisg-nef.streamlit.app/). The platform provides approximately 0.078–2 vCPUs, 0.69–2.7 GB RAM, and up to 50 GB storage per app [25], which suffices for the ETL, aggregation, and visualization workflows. All evaluations in this chapter are reproducible in this controlled environment, *CareNet*'s **source-code is also available** [18].

Datasets. Two complementary packet-capture datasets were utilized, one from Jiang *et al.* [11], which combines operator-side (BRAS) and customer-side (ONU) traces, achieving over 95% application-identification accuracy. This dataset serves as a semantic link between packet-level descriptors and user activities. The second dataset, from Hjelmvik's [9], offers a continuous 40-day capture with realistic, unlabeled household-like traffic. Together they cover content-aware and time-continuous dimensions needed to test *CareNet's* router-centric mapping of network metrics to behavioral indicators.

B. Evaluation Methodology

Traffic is partitioned into fixed windows (Δ=5 min). From each window we compute counts (packets/bytes), transport shares (TCP/UDP/Other), session gaps, and DNS statistics at registrable domain (SLD/eTLD+1) level. Day-level metrics $x_i(d)$ are mapped to memberships $\mu_i(d) \in [0,1]$ using triangular functions $\mu_i(x; \ell_i, m_i, h_i)$ with optional inversion when "less = more evidence". Per-criterion daily likelihoods are then:

$$L_k(d) = \sum_{i \in \mathcal{F}_k} w_i \, \mu_i(d), \quad \sum_i w_i = 1,$$

and a persistence gate (M=14, N=6, θ=0.6) marks criterion k "present" iff at least N of the last M days satisfy $L_k(d) \geq \theta$. We evaluate four criteria in CareNet (C1, C2, C4, and C8) but, due to space, we focus here on C4 (sleep timing/duration) and C8 (attention/indecisiveness).

C. Criteria C4 and C8 — Sleep and Attention Indicators

Configuration. Criteria C4 and C8 represent two complementary, router-visible behaviors: diurnal rhythm (sleep timing and duration) and attention stability (concentration and decisiveness). Table IV lists the configured metrics, weights, and triangular membership limits.

TABLE IV: Configuration of parameters for C4 (Sleep timing/duration) and C8 (Difficulty concentrating/indecisiveness).

C4/C8 Feature	w	lo	mid	hi	Dir.
C4 — Sleep timing / duration					
C4_F2 WakeAfter0400Min	0.65	120	1085	1085	↑
C4_F4 SleepDurationZAbs30d	0.20	0.25	0.80	0.80	↑
C4_F7 DaytimeIdleRatio0818	0.05	0.00	0.08	0.16	↑
C4_F8 NightDayTrafficRatioBytes	0.15	0.20	1.00	1.00	↑
C8 — Difficulty concentrating / indecisiveness					
C8_F2 DNSBurstRatePerHour	0.40	25	61	61	↑
C8_F4 RepeatedQueryRatio60m	0.40	0.80	1.00	1.00	↑
C8_F8 MedianIKSsec	0.20	0.12	0.22	0.22	↑

MF: triangular; Dataset B [ALL_OTHER]; M=14, N=6, θ=0.6, τ=1.

Notes: w = weight; MF = membership function; lo/mid/hi = membership limits; Dir. = direction (↑ = increase indicates evidence).

Influence of features. Figure 5 illustrates how raw packet telemetry is translated into interpretable behavioral likelihoods through the FASL pipeline. The diagrams highlight that each feature operates as a deterministic mapping from measurable network primitives, *e.g.*, packet counts, inter-arrival gaps, or DNS lookups—into bounded evidence values that together characterize either diurnal rhythm (C4) or attention stability (C8).

For C4, the dominant signal is *WakeAfter0400Min* (weight 0.65), derived from the last observed network activity before 04:00 local time. It anchors the detection of delayed

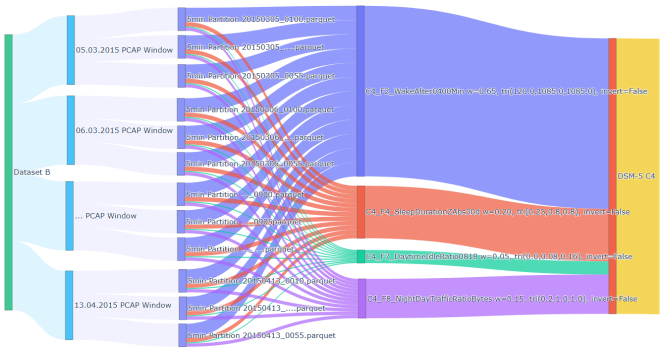

(a) C4 — Sleep timing and duration.

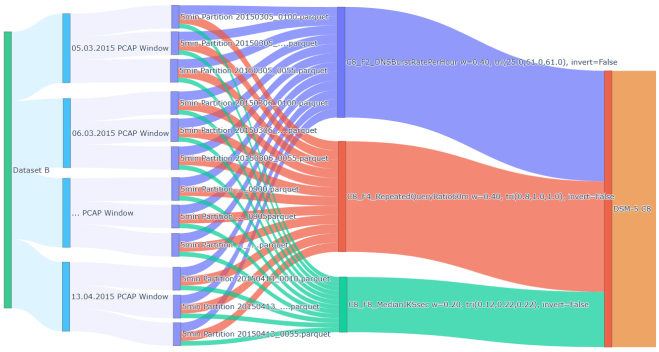

(b) C8 — Difficulty concentrating / indecisiveness.

Fig. 5: Information flow for Criteria C4 and C8 showing feature extraction and aggregation into DSM-5–aligned indicators.

or fragmented sleep patterns by identifying nights where activity persists abnormally late. The secondary features refine this estimate: *SleepDurationZAbs30d* quantifies how much total nightly inactivity deviates from the 30-day baseline, *DaytimeIdleRatio0818* measures the proportion of idle time during the main activity window (08:00–18:00), and *NightDayTrafficRatioBytes* compares nocturnal and daytime byte volumes. Together they approximate a digital circadian profile purely from transport- and session-level metadata, showing that even coarse temporal packet aggregates can encode sleep-wake regularity without inspecting payloads.

For C8, the decisive features are *DNSBurstRatePerHour* and *RepeatedQueryRatio60m* (weights 0.40 each). The first measures how often multiple distinct second-level domains are contacted in rapid succession, indicating short attention spans or topic switching; the second counts repeated DNS queries for identical domains within an hour, which may reflect checking or uncertainty behaviors. Both are observable directly from resolver traffic and require no flow reconstruction. The auxiliary feature *MedianIKSsec* (weight 0.20) estimates inter-keystroke gaps inferred from packet bursts in interactive sessions, providing an orthogonal, human-interaction dimension.

Membership functions. The contribution stacks in Figure 6 visualize how each feature's membership value contributes to the overall daily likelihoods. These plots make the internal mechanics of the FASL fusion explicit: all metrics are first mapped through bounded triangular functions, then linearly combined according to their configured weights (*cf.* Table IV). These resulting traces show how much evidence each metric provides on each day, making the additive nature of the inference directly observable.

For C4, the dominant feature *WakeAfter0400Min* drives most of the accumulated membership, reflecting persistent late-night activity visible at the packet level. The other features—*SleepDurationZAbs30d*, *DaytimeIdleRatio*, and *Night/DayRatio*—contribute short pulses that refine this baseline by accounting for sleep duration irregularities, daytime inactivity, and nocturnal traffic surges. Together they model both persistent and transient disruptions in

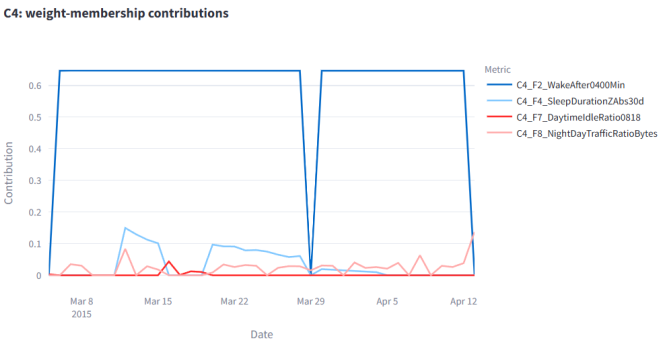

(a) C4 — weight–membership over time.

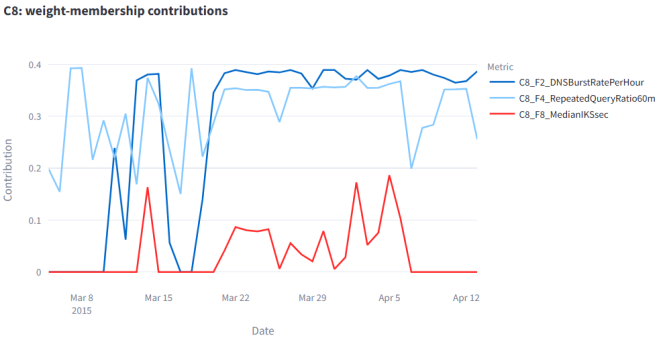

(b) C8 — weight–membership over time.

Fig. 6: Weight–membership contributions for C4 and C8.

digital circadian rhythm. The visual plateau in Figure 6(a) demonstrates that the membership saturates near its upper bound for extended periods, a behavior desirable in fuzzy logic because it prevents single anomalies from dominating the criterion.

For C8, the decomposition in Figure 6(b) shows how attention-related features interact over time. Early in the sequence, peaks in *RepeatedQueryRatio60m* indicate short-lived bursts of redundant DNS lookups, followed by a sustained plateau dominated by *DNSBurstRatePerHour*. This reflects periods of rapid, topic-switching network behavior, as seen in the frequent resolver activity across distinct domains.

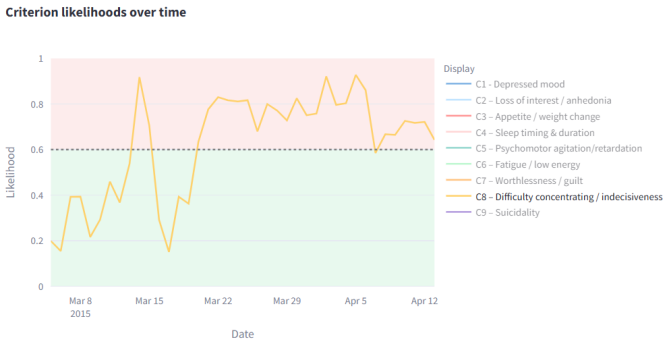

(a) C8 — criterion likelihood over time. Multi-metric alignment produces sustained periods above the gate; isolated fluctuations only nudge the curve.

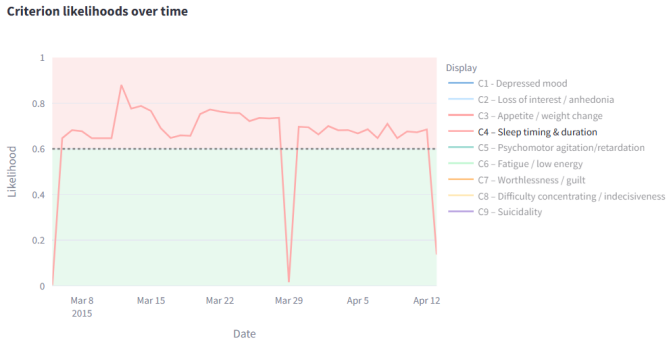

(b) C4 — criterion likelihood over time. Long stretches above the gate reflect a saturated wake-after-04:00 signal with small refinements from the other metrics and a single reset day near the end of March.

Fig. 7: Daily criterion likelihoods for C8 and C4 after FASL aggregation and DSM-style gating.

Criterion likelihood. Figure 7 presents the aggregated daily likelihoods for Criteria C8 and C4 after combining all weighted feature memberships through FASL. Each curve represents the day-level probability that the corresponding behavioral criterion is active, filtered by the temporal gate parameters (M=14, N=6, θ=0.6).

For C8, the likelihood trace shows clear multi-day periods where DNS-related metrics align above the decision threshold, indicating consistent evidence of attentional fragmentation. These plateaus arise when both *DNSBurstRatePerHour* and *RepeatedQueryRatio60m* remain elevated for consecutive days, while intermittent keystroke-gap pulses produce smaller transient increases. This shows the gate acts as a temporal integrator, requiring multiple evidence sources to confirm anomalies within a window.

In contrast, C4 exhibits extended intervals above the gate dominated by the persistent *WakeAfter0400Min* signal, which reflects continuous late-night network activity in Dataset B. Secondary features such as *SleepDurationZAbs30d*, *DaytimeIdleRatio*, and *Night/DayRatio* modulate these plateaus, adding sensitivity to short-term variations in daily rhythm. The stable sections above θ therefore represent corroborated, time-consistent deviations rather than single-day outliers.

D. The User Journey

The final overview aggregates all DSM-5 gate outcomes into a unified, interpretable interface within *CareNet*. It reports three key indicators derived from the same temporal parameters (M=14, N=6, θ=0.6, τ=1): (i) overall episode likelihood, (ii) core-symptom activation, and (iii) the count of criteria currently exceeding their gates within the 14-day window. This stage closes the analytical loop by translating the per-criterion likelihoods into a system-level diagnostic summary.

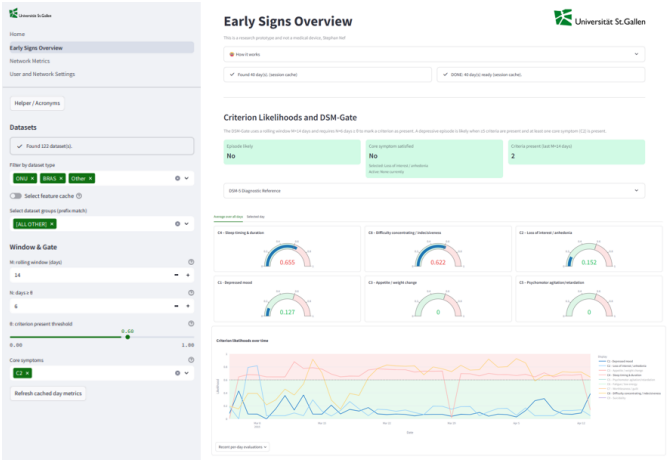

Fig. 8: Entry point *Early Signs Overview*. The DSM-5 gate and triage elements are summarized at a glance for quick interpretation.

Figure 8 introduces this final layer of analysis, where the DSM-style gate is applied globally to produce an at-a-glance behavioral snapshot. In the evaluated configuration, no episode is detected and no core symptom (C1, C2) exceeds its gate, while two non-core criteria, C4 (sleep) and C8 (attention), remain persistently elevated. This outcome reproduces the per-criterion trends observed earlier: stable deviations in circadian and attentional metrics without concurrent changes in mood or motivation. The gate follows DSM-5 logic by requiring multiple corroborating symptoms within a fixed time horizon, translating clinical persistence into a computationally verifiable condition.

This rule-based formulation counts active criteria k with $L_k(d) \geq \theta$ for at least N of the past M days and reports both the instantaneous and aggregated state. Because the configuration and thresholds are explicit, every high-level decision can be traced back to specific metrics and timestamps, ensuring auditability and reproducibility. From a systems viewpoint, this layer acts as the final aggregation of a multi-stage signal-processing chain: packet-level measurements $\rightarrow$ daily features $\rightarrow$ criterion likelihoods $\rightarrow$ interpretable episode-level indicators.

Figures 9 and 10 depict the second layer of interpretation. The gauge panel aggregates mean likelihoods across the observation window, ranking C4 (0.655) and C8 (0.622) highest. The multi-criterion time series confirms this relationship, showing parallel trajectories of sleep and attention features

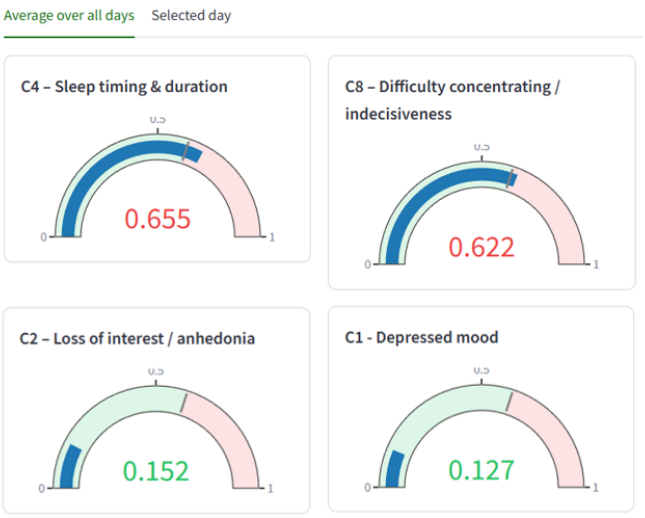

Fig. 9: Sorted gauges overview. Criteria are ranked by average likelihood, highlighting persistent elevations in C4 and C8.

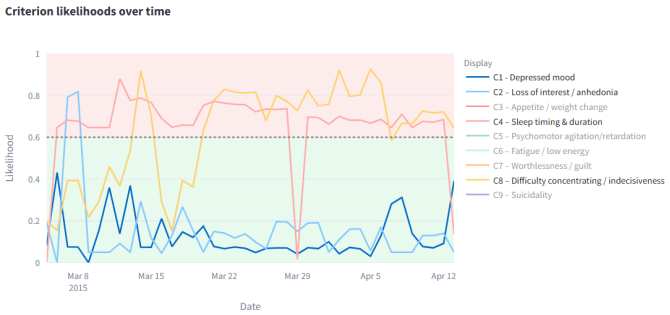

Fig. 10: All criteria over time. Co-movement of C4 and C8 dominates, while C1 and C2 remain near baseline.

with minor, short-lived excursions for mood-related criteria. Together, these visualizations show the scalability of the approach, from per-packet evidence to daily metrics and daily metrics to criterion likelihoods, and finally to explainable depression symptoms via the dashboard.

VI. Discussion and Limitations

The evaluation shows three aspects that matter in practice. First, the method is *explainable by design*. Each metric comes with an explicit baseline (`lo/mid/hi`), a visible membership curve and a time-series of outputs. Users can see how moving a limit or changing a weight changes the indicator. Second, the approach is *person-centric*. Healthy baselines are set per user and outliers above `hi` are intentionally ignored, which reduces noise while keeping unusual days visible in the raw plots. Third, the DSM-style gate aggregates day-level likelihoods in a way that "feels familiar" to clinicians: it requires corroboration across metrics and at least one core symptom before declaring that an episode is likely.

Deterministic, auditable, and open-source. This is a key aspect that makes our output explainable. Every indicator is explained by a visible rule set rather than a black-box model. All packet-level primitives (*e.g.*, inter-arrival structure, direction/volume, DNS activity) map one-to-one to daily features, and triangular memberships convert deviations into bounded evidence; and the persistence gate (M,N,θ,τ) enforces temporal corroboration. Our introduced FASL pipeline behaves like a signal-processing stack down-sampling high-rate traffic into low-frequency behavioral metrics with explicit sensitivity and thresholds.

Privacy considerations. Both network metadata and mental-health indicators are extremely sensitive information. The edge-first architecture ensures that all computation and data storage stay within the household. Only packet headers and derived, day-level features are used, while payloads are never inspected and no data leave the home by default. This approach minimizes risk, aligns with data-minimization principles, and increases user trust.

Limitations: What a router can and cannot observe. A network-centric view can provide several meaningful behavioral metrics: diurnal rhythm (*e.g.*, night/day balance, quiet-hour structure), but not all DSM-5 depression indicators are reliably inferable from network traffic alone (*e.g.*, mood valence, guilt/worthlessness, ideation content, detailed psychomotor signs). In addition, while current likelihoods and weights are currently calibrated against statistical baselines, note that these require clinical/behavioral ground-truth (*e.g.*, PHQ-9, actigraphy [12]). In this regard, *CareNet* allows to tune each weight/parameter in a way that it could be extended to complement existing clinical practice (rather than replacing clinical assessment).

VII. Summary and Future Work

This paper presented *CareNet*, a router-centric framework that transforms household network telemetry into interpretable behavioral indicators aligned with DSM-5 symptom domains. *CareNet* presents a complete edge-based pipeline, from packet capture and feature extraction to fuzzy aggregation and DSM-style gating, showing that clinically relevant temporal patterns can be inferred from header-level data alone.

Future work will focus on validating these indicators against clinical ground truth, extending the framework to additional DSM-5 domains, and automating parameter tuning to adapt thresholds and weights dynamically across heterogeneous household environments.

VIII. Ethical Statement

CareNet is based on Stephan's master thesis at the University of St. Gallen. AI tools were used for spelling, grammar, style, and code assistance in prototype development. The conceptualization, methodological design, and interpretation of results originated from the authors, while AI tools served only as auxiliary instruments for language refinement and implementation support.

References

[1] D. Alamoudi, I. Nabney, and E. Crawley, "Evaluating the effectiveness of the sleeptracker app for detecting anxiety-and depression-related sleep disturbances," *Sensors*, Vol. 24, No. 3, p. 722, 2024.

[2] A. P. Association, *Diagnostic and statistical manual of mental disorders: DSM-5-TR*, fifth edition ed., Washington, DC, 2022.

[3] A.-M. Bucur, "Leveraging llm-generated data for detecting depression symptoms on social media," in *International Conference of the Cross-Language Evaluation Forum for European Languages*. Springer, 2024, pp. 193–204.

[4] F. Cacheda, D. Fernandez, F. J. Novoa, and V. Carneiro, "Early detection of depression: social network analysis and random forest techniques," *Journal of medical Internet research*, Vol. 21, No. 6, p. e12554, 2019.

[5] Z. M. Clayborne, M. Varin, and I. Colman, "Systematic review and meta-analysis: adolescent depression and long-term psychosocial outcomes," *Journal of the American Academy of Child & Adolescent Psychiatry*, Vol. 58, No. 1, pp. 72–79, 2019.

[6] M. De Choudhury, E. Kiciman, M. Dredze, G. Coppersmith, and M. Kumar, "Discovering shifts to suicidal ideation from mental health content in social media," in *Proceedings of the 2016 CHI conference on human factors in computing systems*, 2016, pp. 2098–2110.

[7] Docker Inc., "Docker," 2025, used for containerization and deployment of applications. [Online]. Available: https://www.docker.com/

[8] L. Hale and S. Guan, "Screen time and sleep among school-aged children and adolescents: a systematic literature review," *Sleep medicine reviews*, Vol. 21, pp. 50–58, 2015.

[9] E. Hjelmvik and S. A. F. CERT, "Hands-on network forensics," Workshop material, 2015, fIRST 2015 Berlin, including 40-day continuous capture dataset. [Online]. Available: https://share.netresec.com/s/NBbonFnsZ4HqCTW?dir=/&editing=false&openfile=true

[10] IAB Tech Lab, "IAB Tech Lab content taxonomy v2.2," https://iabtechlab.com/standards/iab-content-taxonomy/, 2023.

[11] W. Jiang, B. Zhang, Q. Zhu, C. Liao, and W. Wang, "A real network environment dataset for traffic analysis," *Scientific Data*, Vol. 12, No. 1, p. 756, 2025.

[12] K. Kroenke, R. L. Spitzer, and J. B. Williams, "The phq-9: validity of a brief depression severity measure," *Journal of general internal medicine*, Vol. 16, No. 9, pp. 606–613, 2001.

[13] K. Kushlev, J. Proulx, and E. W. Dunn, ""silence your phones" smartphone notifications increase inattention and hyperactivity symptoms," in *Proceedings of the 2016 CHI conference on human factors in computing systems*, 2016, pp. 1011–1020.

[14] P. M. Mammen, C. Zakaria, T. Molom-Ochir, A. Trivedi, P. Shenoy, and R. Balan, "Wisleep: Inferring sleep duration at scale using passive wifi sensing," *arXiv preprint arXiv:2102.03690*, 2021.

[15] P. D. McGorry and C. Mei, "Early intervention in youth mental health: progress and future directions," *BMJ Ment Health*, Vol. 21, No. 4, pp. 182–184, 2018.

[16] Mozilla Foundation, "Public suffix list: Documentation," https://publicsuffix.org/, 2024.

[17] K. O. E. Müller, D. Datsomor, D. Schumm, B. Rodrigues, and B. Stiller, "Big brother is watching you: Non-intrusive zigbee user profiling," in *Proceedings of the 20th International Conference on Network and Service Management (CNSM)*. IFIP, 2024. [Online]. Available: https://www.alexandria.unisg.ch/server/api/core/bitstreams/09d68240-b1e3-4481-8d45-c766ed7f5f82/content

[18] S. Nef, "app_pcap_to_dsm5_public: Pcap → dsm-5 streamlit prototype," https://github.com/nefste/app_pcap_to_dsm5_public, 2025, version-controlled research codebase; includes Streamlit app, metrics, ETL, and Docker/Compose assets.

[19] Y. Oguchi, A. Nakagawa, M. Sado, and et al., "Potential predictors of delay in initial treatment contact after the first onset of depression in japan: A clinical sample study," *International Journal of Mental Health Systems*, Vol. 8, p. 50, 2014. [Online]. Available: https://link.springer.com/article/10.1186/1752-4458-8-50

[20] S. H. Pedersen, K. B. Stage, A. Bertelsen, P. Grinsted, P. Kragh-Sørensen, and T. Sørensen, "Icd-10 criteria for depression in general practice," *Journal of affective disorders*, Vol. 65, No. 2, pp. 191–194, 2001.

[21] S. Saeb, T. R. Cybulski, S. M. Schueller, K. P. Kording, and D. C. Mohr, "Scalable passive sleep monitoring using mobile phones: opportunities and obstacles," *Journal of medical Internet research*, Vol. 19, No. 4, p. e118, 2017.

[22] S. Saeb, M. Zhang, C. J. Karr, S. M. Schueller, M. E. Corden, K. P. Kording, D. C. Mohr *et al.*, "Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: an exploratory study," *Journal of medical Internet research*, Vol. 17, No. 7, p. e4273, 2015.

[23] J. Smith-Nielsen, S. Matthey, T. Lange, and M. S. Væver, "Validation of the edinburgh postnatal depression scale against both dsm-5 and icd-10 diagnostic criteria for depression," *BMC psychiatry*, Vol. 18, pp. 1–12, 2018.

[24] C. Stachl, Q. Au, R. Schoedel, S. D. Gosling, G. M. Harari, D. Buschek, S. T. Völkel, T. Schuwerk, M. Oldemeier, T. Ullmann *et al.*, "Predicting personality from patterns of behavior collected with smartphones," *Proceedings of the National Academy of Sciences*, Vol. 117, No. 30, pp. 17 680–17 687, 2020.

[25] Streamlit Inc., "Manage your app — streamlit community cloud," 2025, accessed: 2025-10-18. [Online]. Available: https://docs.streamlit.io/deploy/streamlit-community-cloud/manage-your-app

[26] F. Svenaeus, "Diagnosing mental disorders and saving the normal: American psychiatric association, 2013. diagnostic and statistical manual of mental disorders, american psychiatric publishing: Washington, dc. 991 pp., isbn: 978-0890425558. price: $122.70," *Medicine, Health Care and Philosophy*, Vol. 17, pp. 241–244, 2014.

[27] R. Wang, F. Chen, Z. Chen, T. Li, G. Harari, S. Tignor, X. Zhou, D. Ben-Zeev, and A. T. Campbell, "Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones," in *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*, 2014, pp. 3–14.

[28] R. Wang, W. Wang, A. DaSilva, J. F. Huckins, W. M. Kelley, T. F. Heatherton, and A. T. Campbell, "Tracking depression dynamics in college students using mobile phone and wearable sensing," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, Vol. 2, No. 1, pp. 1–26, 2018.

[29] S. Ware, C. Yue, R. Morillo, J. Lu, C. Shang, J. Bi, J. Kamath, A. Russell, A. Bamis, and B. Wang, "Predicting depressive symptoms using smartphone data," *Smart Health*, Vol. 15, p. 100093, 2020.

[30] S. Ware, C. Yue, R. Morillo, J. Lu, C. Shang, J. Kamath, A. Bamis, J. Bi, A. Russell, and B. Wang, "Large-scale automatic depression screening using meta-data from wifi infrastructure," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, Vol. 2, No. 4, pp. 1–27, 2018.

Appendix

Repository Access and Reproducibility

Source code and live deployment. All components required to reproduce the results presented are publicly available. The following notes describe where to find the source code, how to run the application locally, and how the repository is structured, and build the system in Figure 11.

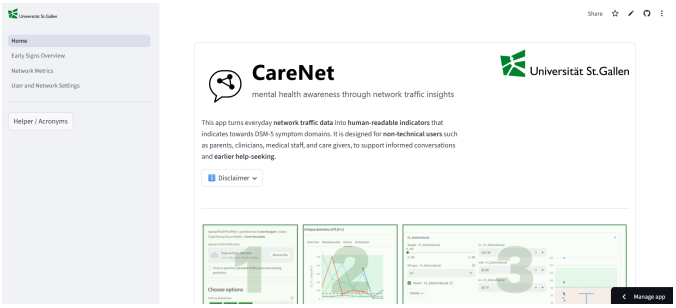

Fig. 11: CareNet's landing page.

The full implementation of *CareNet* is hosted on GitHub at `nefste/app_pcap_to_dsm5_public`. It includes the complete Streamlit dashboard, metric definitions, ETL utilities, and configuration files used in the evaluation. A live instance is available at https://unisg-nef.streamlit.app/. Access credentials can be obtained directly from the author upon request.

Running locally. In addition, *CareNet* can be executed either through Docker or within a Python virtual environment. The following commands reproduce the full setup:

```
cd app
docker compose up --build
# Application available at http://localhost
    :8501
```

Listing 1: Run locally using Docker

```
cd app
python -m venv ..\.venv
..\.venv\Scripts\activate    # on Windows (
    adjust for macOS/Linux)
pip install -r requirements.txt
streamlit run 00_Home.py
```

Listing 2: Run locally using a Python virtual environment

If additional utilities are required, install repository-level dependencies using `pip install -r ../requirements.txt`. When `libpcap` is unavailable, set `SCAPY_USE_LIBPCAP=no`; the application reads PCAP files from disk and does not require live packet capture.

Repository structure. The main files and directories of the application are summarized below. Each entry lists the purpose of the corresponding component and its role in the overall system.

- **A1** `app/00_Home.py` — Landing page with disclaimer and general overview of the dashboard.
- **A2** `app/pages/01_Early_Signs_Overview.py` — Loads and normalizes Parquet data; computes and visualizes FASL and DSM-Gate outputs.
- **A3** `app/pages/02_Network_Metrics.py` — Defines per-criterion metric computations and corresponding UI views.
- **A4** `app/pages/03_User_and_Network_Settings.py` — Manages profiles, IP mappings, uploads, and ETL partitioning.
- **A5** `app/metrics/common.py` — Contains helper functions for sessionization (`sessions_from_timestamps`), domain parsing (eTLD+1), and data formatting.
- **A6** `app/metrics/base_features.py` — Generates daily base features such as *LateNightShare*, *LongestInactivityHours*, and IS/IV ratios.
- **A7** `app/metrics/criterion1.py-criterion9.py` — Implements DSM-5 criterion-specific metrics and their respective likelihood computations.
- **A8** `app/auto_tune.py` — Optional module for auto-tuning triangular memberships and feature weights.
- **A9** `app/processed_parquet/<dataset>/` — Stores five-minute Parquet files linking ETL and analytics stages (`<dataset>__YYYYMMDD_HHMM.parquet`).

Loading Traces and Configuring User Profiles

CareNet supports interactive exploration of packet-capture datasets directly within the dashboard.

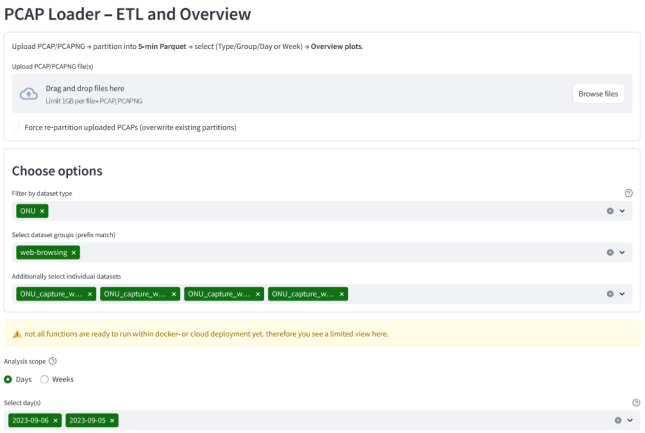

Fig. 12: Configuration page for selecting datasets and analysis periods in the *CareNet* dashboard.

After launching the application, users can open the *User and Network Settings* page to configure which PCAPs to load and how they are partitioned (*cf.* Figure 12). The interface allows selecting one or more dataset groups such as `BRAS`, `ONU`, or `snort.log`, as well as thematic subsets (e.g., *video*, *web-browsing*, or *gaming*).

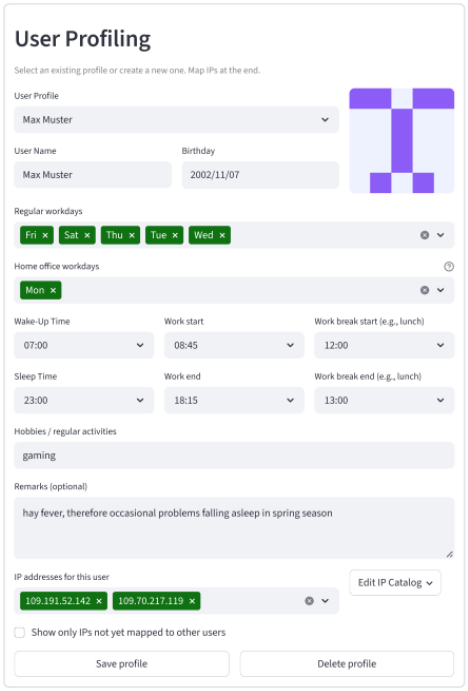

Fig. 13: User profiling and IP mapping used by direction-aware metrics and interpretations.

It is also possible to maintain profiles (wake/sleep habits, workdays, notes) and map local IPs to users (Figure 13). This mapping activates lightweight direction heuristics (private→public ≈ upstream, public→private ≈ downstream), improving "active vs. passive" features without content access.