A copula approach to modelling dependency patterns in routine vital sign observation timeliness

Rupert Ironside-Smith
School of Computer Science and Informatics
Cardiff University
Cardiff, United Kingdom
Ironside-SmithR@Cardiff.ac.uk

Stuart Allen
School of Computer Science and Informatics
Cardiff University
Cardiff, United Kingdom
AllenSM@Cardiff.ac.uk

Beryl Noë
School of Computer Science and Informatics
Cardiff University
Cardiff, United Kingdom
noeb@Cardiff.ac.uk

Liam Turner
School of Computer Science and Informatics
Cardiff University
Cardiff, United Kingdom
TurnerL9@Cardiff.ac.uk

***Abstract*—Healthcare staff in a hospital ward setting typically monitor patients by taking vital sign observations at regular intervals, usually every 6 - 12 hours for routine observations, but more frequently for critical patients. Patients on similar schedules have been shown to be regularly batched in a practice called 'ward rounds', but to what extent healthcare staff manage observations on 'non-routine' intervals independently to those scheduled for the subsequent ward round has yet to be established. This study examines the Time-To-Next-Observation (TTNO) for vital sign observations with planned schedules (e.g., 1 hour) as random variables defined by their hazard distribution functions. Joint distribution functions sampled from any pair of TTNO distributions could be used to calculate the probability that any two observations on set schedules will happen collectively. However, it is clear that this model fails to capture underlying dependency structures seen in empirical results. We propose a copula approach to extract and quantify pairwise non-linear relationships for all standard observation intervals across 20 study wards. This study showed that most wards operate with significant levels of dependency between observation schedules, largely reflecting broad ward characteristics referenced in other works, yet present a deeper level of insight into individual ward operations. Understanding current levels of dependency between regular observation scheduling and routine operations has the potential to become essential knowledge for ward stakeholders when designing staff resource strategies.**

***Index Terms*—Hospital Ward, Vital Signs Observations, Secondary Data Analysis, Hazard Function, Copula Method**

I. Introduction

In a typical hospital ward setting, healthcare staff routinely complete vital sign observations on all patients. Within the Welsh National Health Service (NHS Wales) for example, this consists of 6 measurements in accordance to the *National Early Warning Score 2* (NEWS-2) standard; blood pressure, heart rate, respiratory rate, temperature, level of consciousness, and oxygen saturation. The NEWS-2 standard provides a simple and effective method of aggregating the individual vital sign values into a single digit score which can be used to calculate the target gap to a patient's subsequent observation, known as the *observation interval* (I). Patients on a ward are typically observed at regular intervals (e.g., two, three, or four times a day [12], [14], [21]), subject to the requirements of health monitoring policies and the last reported vital signs, but individual patients may need additional observations. The combination of individual patient and routine observation intervals creates a notably non-uniform daily pattern of vital signs observations within a ward (that is, the hourly volume of vital sign observations changes throughout the day) [12], [16], with clinical staff typically consolidating most routine patient observations into 'ward rounds' [1]. When patients require individual plans of additional observations, these may also be either pulled forward or pushed back to align with these rounds. These patterns reflect operational management at the ward level, including how staff implement policies, and conduct observations around other activities such as visitation times, meals, medications, and sleep. Stratifying vital signs observation timeliness with respect to the planned Time-To-Next-Observation (TTNO) has shown that patients with shorter planned intervals and high Early Warning Scores (EWS)[1] tend to have the most late or missed observations [14], [23], [24]. However, the specific causes behind late observations have not been widely studied, and existing works rarely mention observations that are pulled forward ahead of schedule. Additionally, there has been little consideration of how timeliness relates to planned observation intervals. In this paper, we start looking at the regular interactions between observations on standard schedules, such as how often observations on a 1 hour schedule coincide with those on 8 hour schedules, and how these compare to current ward-level metrics like timeliness

[1]In a busy ward environment it is slow for healthcare staff to concisely describe the severity of a patient case using multiple individual parameters (e.g., blood pressure, heart rate, respiratory rate, temperature, level of consciousness, and oxygen saturation). So, staff combine multiple vital sign measurements and clinical judgement into a convenient single-digit number using a standardised scoring system, called an *Early Warning Score* [7], [17].

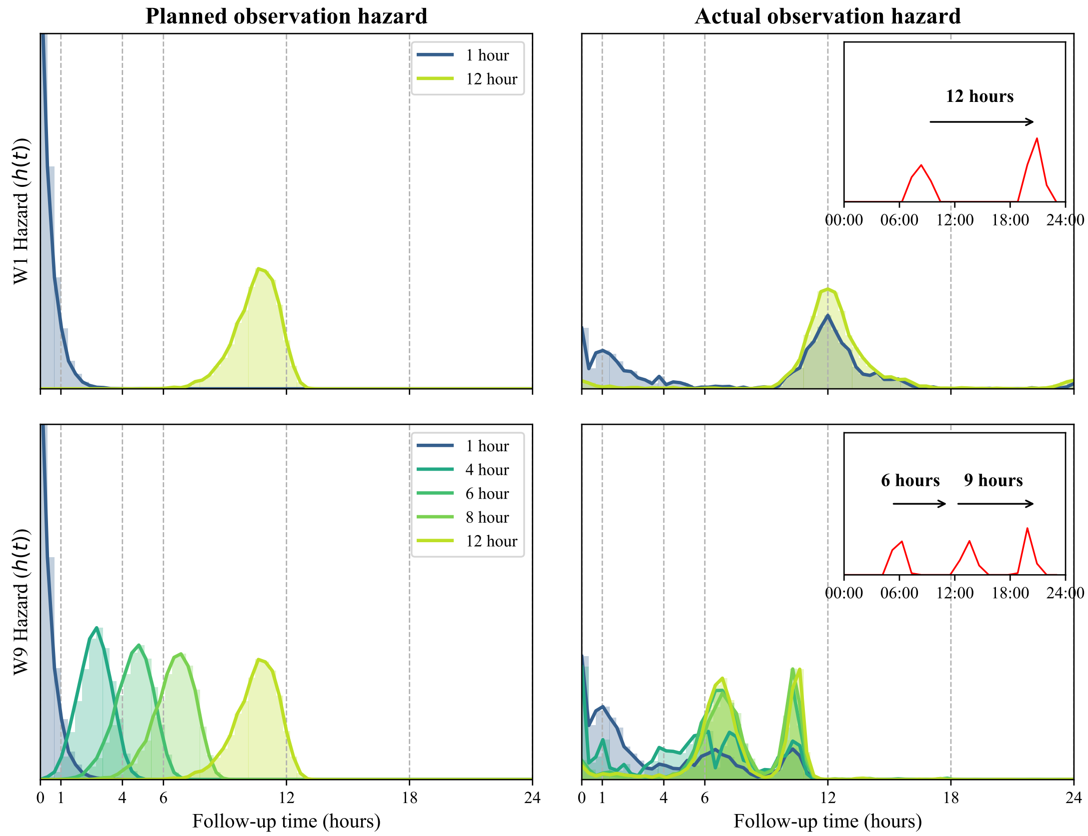

Fig. 1. The difference between planned and empirically derived observation hazard distributions ($h(t)$) for a range of observation intervals (1 hour, 4 hour, 6 hour, 8 hour, and 12 hour) across two example wards, W1 (top) and W9 (bottom). Inset axes dictate daily observation volume patterns in each ward and annotated with the average time gap between routine ward rounds.

conformance.

A. *Problem definition*

If individual vital sign observations within a dataset can be represented with the following features:

$$\text{Observation ID } (i),\ \text{ward ID, patient ID } (p),\ \text{time } (t),\ I \quad (1)$$

The dataset can be easily transformed into time-to-event recordings by calculating the TTNO of all observation in the data as the actual time differential be two subsequent observations on the same patient [12]. In cases where there is not a subsequent observation, or the TTNO exceeds 24 hours (e.g., patients between visits [5]), the outcome is considered 'right-censored' in line with convention [13]. Right-censoring data when using electronic health record datasets can lead to structural incompleteness [3], but for vital sign observations this can be mitigated by setting the TTNO as 24 hours, the end of our follow-up period. With this data, empirical Nelson-Aalen hazard distributions [4] can be derived to describe the probability that an observation with a specified schedule (i.e., the observation interval) will be taken within its interval. Using this model, the TTNO for a patient can be defined as a random variable, T, with its probability distribution of occurring at time t defined by its hazard function, $h(t)$.

Figure 1 illustrates the current discrepancy between planned hazard distributions for different observation intervals, which we model as Weibull distributions using shape (k) and scale (λ) parameters proportional to the observation interval (I) (described in Equation 2)[2], and the equivalent empirically derived hazard distributions in two example wards. These two wards were chosen as they are representative of common observation interval distribution groupings seen in this dataset and other work [21]. Out of the 20 wards in this study, wards W1-W8 are in 'group 1', a low-severity ward that predominantly uses longer intervals (i.e., likely maintains typically stable patients), and wards W9-W20 are in 'group 2', utilise a more diverse range of observation schedules.

$$k = I \text{ (in hours)}, \quad \lambda = k - \frac{k}{k + \exp\left(\frac{1-k}{2}\right)} \qquad (2)$$

Although the assumption that patient care is individualised with respect to vital sign observations (i.e., the stochastic

[2]What the hazard distribution for a *so far as is reasonably practicable* management of a given observation interval *should* look like is currently unknown, however, preliminary analysis suggests a distribution that accommodates potential variations in timing (such as overseeing a student observation, or a patient away from their bed). But, this is proposed solely as a baseline comparative metric for current ward management practices, not a target for hospital ward stakeholders to optimise or strictly adhere to.

behaviour of one patient's TTNO is fully described by the survival function of the observation interval and patient condition) is appropriate when considering the individual effects of patient condition on the TTNO, one must think of what underlying effects of timeliness dependence between patients during routine ward operations exist. Suppose two patients in neighbouring beds both have their vital signs observed during the same ward round and both present no symptoms of onset deterioration, it is not a stretch to imagine that their subsequent observations would both then occur in sequence during the next routinely scheduled ward round (i.e., their TTNOs are largely dependent). What if their planned observation intervals were different, or if one patient was showing signs of onset deterioration? Measuring the level of vital sign observation timeliness dependency in wards with limited staffing resources will unveil a novel aspect of ward operations, and by which we approach by answering the following research questions:

1 Can the dependent nature of the TTNO between different vital sign observation schedules be appropriately characterised? And if so, which observation schedule pairings (e.g., 6 hours and 8 hours, or 1 hour and 12 hours) are most dependent during routine ward operations?

2 Are inter-schedule dependencies consistent across wards? If not, are there patterns which align to ward-level characteristic groupings seen in other studies?

B. Data overview

This study works with a large, anonymised vital signs observation dataset ($N_{obs} = 770,720$) compiled from 20 wards across 7 hospital sites and 8 specialisms (Medical, Surgical, Rehabilitation, Care Of The Elderly, Orthopaedic, Cardiology, and Acute Stroke) run by a health board in Wales, UK. For this study, for each vital signs observation we consider the features described in Equation 1, but other information (e.g., NEWS or staff ID) may be included in future studies [11], [21].

II. Method

To determine whether the TTNO for a patient on a selected observation interval, which we characterise by the hazard distribution T_U, is dependent on the TTNO of another patient with a specific interval, T_V, we can evaluate the correlative features seen in the joint distribution function $F(T_U, T_V)$.

A. Estimating marginals

As mentioned in Section I-A, we can estimate the hazard function for the TTNO of a patient with a selected observation interval by using Nelson-Aalen methodology [4] (shown for our example wards in Figure 1). Simulating a joint distribution function from any pair of marginals, $F(T_U, T_V)$, can therefore provide an estimate to the probability that T_U and T_V would occur concurrently[3]. However, Figures 2 and 3 visualise what

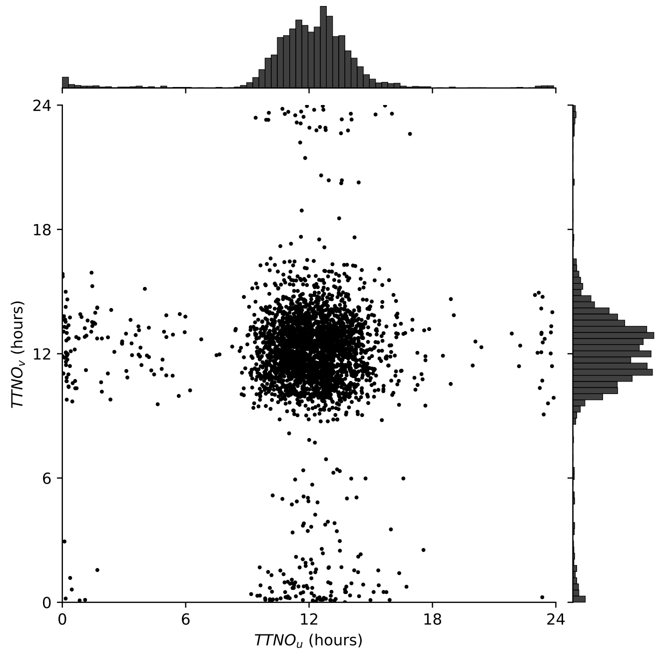

Fig. 2. Simulated joint plot $F(T_{12}, T_{12})$ ($N_{obs} = 3000$). Marginal distributions were derived from Nelson-Aalen hazard estimations for observations with $I = 12$ hours in example ward W1.

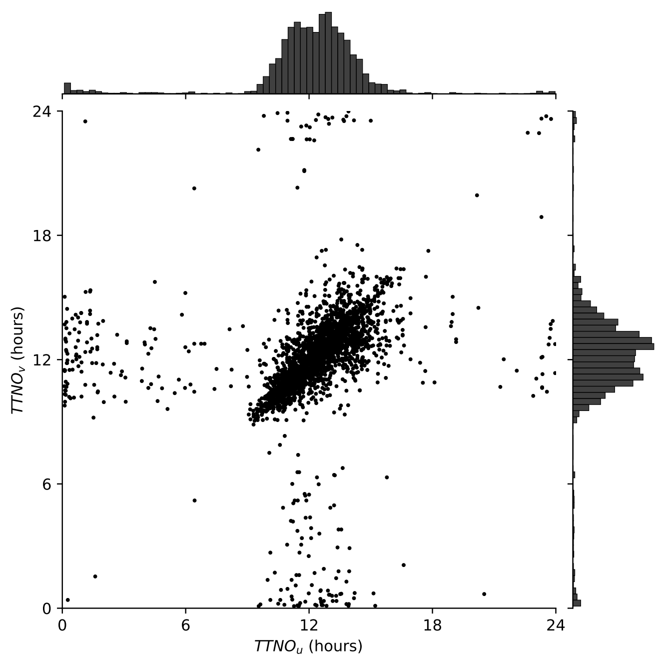

Fig. 3. Empirical joint plot $F(T_{12}, T_{12})$ ($N_{obs} = 3000$) for concurrent observations in ward W1.

discrepancy between simulated joint plots and real data examples may exist, despite being sampled from the same marginal distributions. Simply estimating marginals alone for a pair of observation cases may fail to capture essential dependency structures in vital sign observations management, therefore we

[3]By *concurrently*, we mean within a short timeframe, as clinical staff cannot perform two observations at exactly the same time.

must consider methods to quantify these interactions seen in the data.

B. The copula function

Copula functions were introduced in the context of probabilistic metric spaces as non-linear standardisations of multivariate distributions, and have been commonly used as a transparent method to insert dependency between two known marginal distributions for survival analysis [22]. At its base level, a copula can be described through Sklar's theorem [26] as follows; suppose $X_1, \ldots, X_d$ are random variables with continuous distribution functions $F_1, \ldots, F_d$ and joint distribution function F, then there exists a unique copula C (a distribution function on $[0,1]^d$ with uniform marginals) such that for all $x = (x_1, \ldots, x_d)^\top \in \mathbb{R}^d$:

$$F(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)) \tag{3}$$

Conversely, given any distribution functions $F_1, \ldots, F_d$ and copula C, F defined through equation 3 is a d–variate distribution function with marginal distribution functions $F_1, \ldots, F_d$. Equation 3 can be rewritten for $u = (u_1, \ldots, u_d)^\top \in [0,1]^d$ as:

$$C(u_1, \ldots, u_d) = F(F_1^1(u_1), \ldots, F_d^1(u_d)) \tag{4}$$

Reading Equation 4 from right to left yields the construction of the copula C from any joint distribution function with continuous marginal distribution functions $F_1, \ldots, F_d$. Copula functions come in many variations, but typically appear as one of several common families; Gaussian, Frank, Gumbel, Student-t, and Clayton. Each family inserts specific structures, such as tail dependencies, and are scaled by certain hyper-parameters (for each of the named it is just 1 hyper-parameter, ρ for the Gaussian copula, and θ for the others) [19].

C. Application of the copula function

The copula methodology was popularised through the field of quantitative risk management of financial and insurance products [8], [15]. Despite criticisms that early models were oversimplified and inherently static [6], the copula approach has demonstrated positive results across various other fields, including clinical trials [25], [28], reliability engineering [2], [10], and environmental studies [18], [27]. The transparent and adaptable nature of the copula approach for modelling complex dependence structures without imposing restrictive assumptions on marginal distributions has already proven to be useful across many fields and offers a strong precedent for its use in our application.

D. Copula evaluation and selection

The timing and structure of ward rounds has already been identified by several study groups [12], [14], [21] using the 'activity peaks' in hourly vital sign observation volume curves, like the ones shown in the inset plots in Figure 1. The average time between ward rounds ($\bar{d}$) and the typical length of ward rounds ($\bar{w}$) can be used to identify the most aligned observation interval to the practised ward rounds[4]. As we identify relationships between different observation schedules during routine ward operations, we give some precedence to schedule pairs that coincide with the subsequent ward round (e.g., if ward rounds occur every 12 hours, the dependency between observations on 12 hour intervals will be most reflective of ward round policy influence). It may also be challenging to discern whether short-interval observations taken 'on time' (i.e., an observation that is taken within its specified interval) and in quick succession was due to dependence or simply the tight timeliness bands for these schedules.

Equation 4 can be used to transform the empirical joint distribution function F(T_U, T_V) for any given pair of observation schedules into the underlying copula function, granted that the chosen schedules have been used appropriately frequently ($N_{obs} > 3000$). The derived copulas for each frequent observation schedule tuple can then be fitted against the common Archimedean copula families mentioned at the end of Section II-B. The Log Likelihood measure between the empirically derived copula C(T_U, T_V) and a simulated copula of the same hyper-parameter can be used to indicate the goodness of fit[5].

E. Ward-level comparisons

With an appropriate copula chosen to represent the dependency structure between any given pair of observation schedules, it becomes possible to make pairwise comparisons within a ward and with equivalent pairs in other wards. This analysis aims to provide a literature reference for what constitutes *high* dependence in observation schedule relationships, and determine what ties exist to categories defined from existing work, such as number peaks, observation schedule volume distributions, ward size, and specialism.

III. Results

For clarity, we begin the results section by discussing the workflow and findings using the example ward W1 described in Figure 1, after this, we present copula fitting results for both example wards. These preliminary results will illustrate how the features seen in empirical vital sign observation survival curves align with significant copula structures and indicate what to expect when extending our method across all study wards. The entire study ward set will not only substantiate existing ward-level categorisations seen in the literature but also highlight new nuances and sub categories in routine vital sign observation management.

A. Ward W1 case study

Figure 4 illustrates the Nelson-Aalen hazard distributions for the three observation interval pairings seen within ward W1; (1 hour, 12 hour), (1 hour, 1 hour), and (12 hour, 12

[4]The Scipy Signal Python3 library (Parameters; relative height=0.05, distance=2, prominence=0.05) was used be used to calculate $\bar{d}$ and $\bar{w}$.

[5]The Python3 Copulae library was used to determine the best fit parameters and log-likelihood scores for this data in this study.

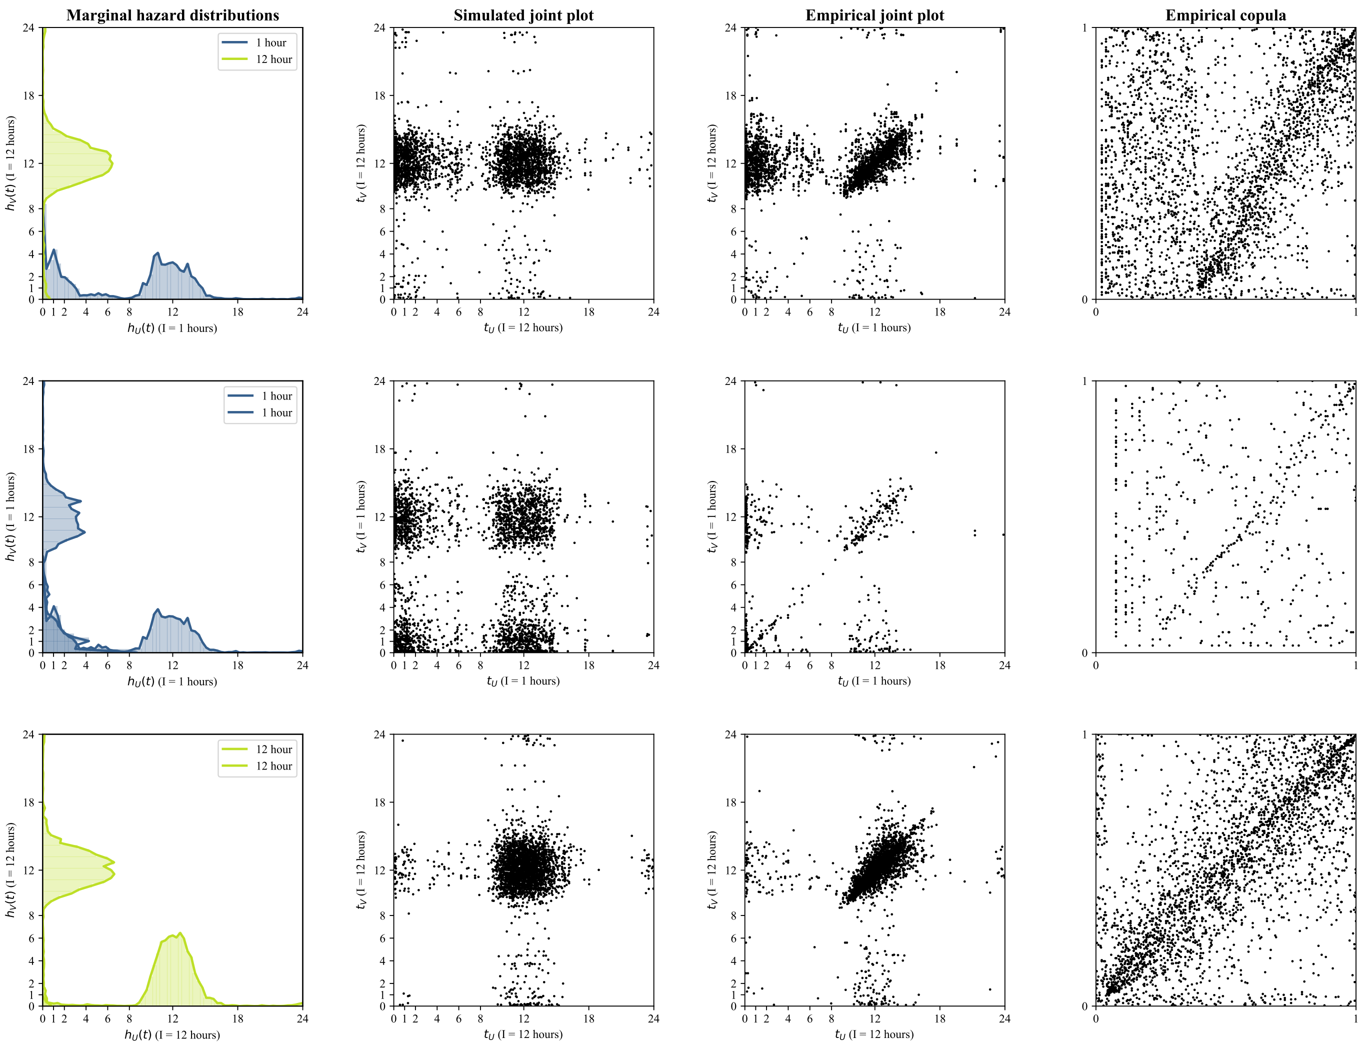

Fig. 4. Left: Nelson-Aalen hazard distributions for the three possible observation interval pairings in ward W1. Centre left: Simulated joint plots for $F(T_U, T_V)$. Centre right: Empirical joint plots for which TTNO values for $F(T_V)$ come from observations which occurred within 15 minutes of observations $F(T_U)$. Right: Derived copula functions, $C(T_U, T_V)$, from the empirical $F(T_U, T_V)$.

hour). As mentioned in Section II-A, this visually reiterates how simply modelling the probability the TTNO outcomes for two observations taken in quick succession based on empirical Nelson-Aalen hazard distributions is insufficient for capturing the real dependent nature of vital sign observations. Looking to the empirical joint plots and the extracted copulas, we see several vital sign observations management structures surface, such as the regularity of which observations on a 1 hour interval are taken around 12 hours, the non-negligible volume of observations on a 12 hour interval taken within 1 hour. In this example ward $\bar{d}$ was calculated as 12 hours, which can be visually corroborated by the inset axis in Figure 1. This reflects in the strong dependency structure in the extracted copula for $F(T_{12}, T_{12})$ and the extracted copula $C(T_1, T_{12})$ also shows this structure, however, only when T_U exceeds 9-10 hours (i.e., when the TTNO aligns with the subsequent ward round).

B. Copula fit to example wards

Tables I and II present calculated copula parameters and the corresponding log-likelihood score (broadly indicating the goodness-of-fit) for Gaussian, Clayton, Frank, Student-t, and Gumbel copula families. For both example wards, the Gumbel copula appeared to be best fitting when the dependency includes observations on a 1 hour interval, and the Student-t copula had highest log-likelihood scores for tuples of longer intervals, but also saw repeated convergence errors. Overall, the Frank [9] copula is the most suitable fit for our data since most ward round schedules align to longer intervals (such as $F(T_{12}, T_{12})$ in Figure 1). This makes it the best representation of the actual dependencies resulting from ward round practices. Several observation schedule pairings had theta parameters exceeding 5 with high corresponding log-likelihoods (see Figure 5 for context), indicating strong dependency structures. Although the Frank copula is a good choice for this model, other copulas fit some pairings well. For instance, there is noticeable lower tail dependency bias between pairs of $I = 12$ hour observations in the empirical joint plot for W1 (see Figure 4). This may be essential feature in specific modelling applications and better represented by using a different copula, such as Clayton copula.

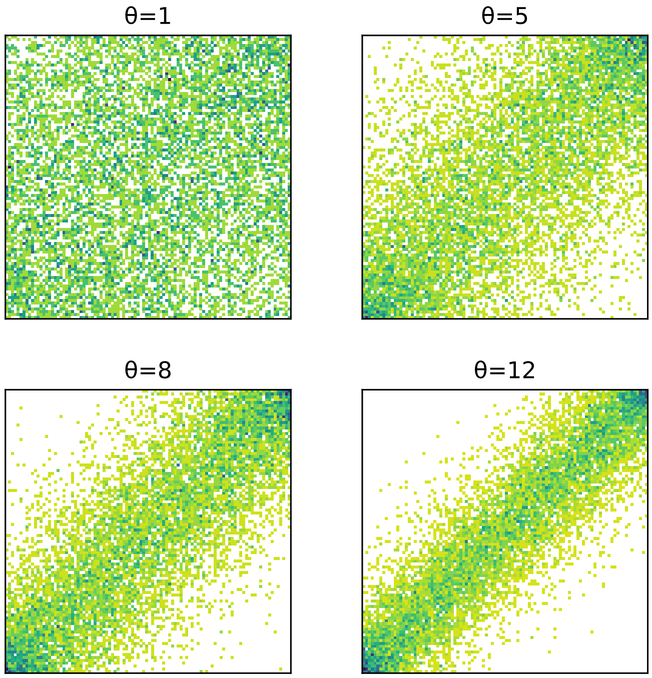

Fig. 5. Example θ parameters (1, 5, 8, and 12) for Frank copula.

C. All study wards

Assuming the Frank copula provides the best overall fit for ward TTNO survival dependency models (for consistency), we can illustrate the dependency structure of any given ward using a simple network diagram. Network edges are weighted by fitted θ parameters and are only visible if they exceed a set log-likelihood threshold. Figure 6 illustrates these networks for our 20 study wards with the threshold log-likelihood of 50 (determined empirically by the results in Table I). The order has been organised with respect to the observation interval distribution (discussed in Section I-A), but with the last three wards in *group 2* (W18, W19, and W20) considered as their own sub-group because of their significantly different daily ward round structures. Using Figure 6, we can make some general comments about the treatment of commonly seen observation intervals;

Short intervals It is unsurprising that with overall fewer utilised intervals in group 1 than in group 2 short-interval observations are commonly paired with 12-hour intervals ($\mu_\theta = 3.847$, $\sigma_\theta = 2.749$), which aligns to the next ward round - But, the same occurs in ward group 2 with a more diverse range of interval lengths, and to the same extent ($\mu_\theta = 3.625$, $\sigma_\theta = 2.713$). This is even higher if excluding the subgroup [W18, W19, W20] ($\mu_\theta = 4.833$, $\sigma_\theta = 1.994$). In ward group 2 (excluding the subgroup), short intervals also show consistent dependency to mid-long-length intervals (to 6 hour intervals; $\mu_\theta = 5.064$, $\sigma_\theta = 1.267$, and to 8 hour intervals; $\mu_\theta = 5.349$, $\sigma_\theta = 1.297$). Conversely, it appears that 1 hour intervals are rarely paired with 4 hour intervals (only in W15), or even with themselves (only in W17).

Mid-length intervals Although used infrequently across ward group 1, they show dependency to longer intervals when they do occur. In group 2, these intervals exhibit higher dependencies between among themselves (6 and 8 hours; $\mu_\theta = 5.978$, $\sigma_\theta = 1.256$) whilst also showing dependencies on long 12 hour intervals (6 and 12 hours; $\mu_\theta = 3.084$, $\sigma_\theta = 3.115$, and 8 and 12 hours; $\mu_\theta = 3.400$, $\sigma_\theta = 3.268$). The periods between ward rounds in this group (between 6-8 hours) likely prompts completing all mid-length and longer intervals together.

Long interval 12 hour intervals are largely dependent on other 12 hour observations in both ward groups (group 1; $\mu_\theta = 4.378$, $\sigma_\theta = 2.537$, and group 2; $\mu_\theta = 6.623$, $\sigma_\theta = 1.766$), although the relationship is more pronounced in ward group 2 despite shorter gaps between rounds and a more diverse range of intervals in use. The central position of 12 hour intervals in the group 1 networks may indicate directionality in the dependence structure, however the small sample size makes it hard to test. Whether the pattern is consistent in group 2 is also not clear, which could be because other intervals that align closer to the ward round gaps have more influence.

When comparing both ward groupings, including the sub-group within ward group 2, we see that observation schedule dependency structures are somewhat aligned to the ward's daily rounds policy. This is corroborated by a statistically significant moderate positive correlation between number of ward rounds and the average θ across the edges present in the network (Spearman's rank correlation $= 0.515$, $p < 0.005$). This reveals a potential self-perpetuating dynamic where staff operate with dependencies dictated by the constraints that they operate within. For example, if a ward's technical setup or policy only allows for routine 12 hour intervals and non-routine 1 hour intervals, it is unsurprising that these intervals are consolidated in cases where staff are under-resourced or if many patients deteriorate concurrently. Conversely, in a busy ward with six or more observation interval options of varying utilisation these results show that appropriately balanced staffing resources is crucial in managing individual patient schedules without continuously overstaffing the ward.

IV. Discussion

Section III successfully used a copula approach to confirm and quantify the presence of non-linear relationships between different vital sign observation schedules during routine ward rounds. The results showed that pairwise schedule dependencies within a ward may be dictated by higher-level ward round

TABLE I

EMPIRICAL COPULA FITTING FOR SCHEDULE PAIRS IN W1. HIGHER VALUES FOR LOG LIKELIHOOD INDICATE A BETTER FIT TO THE COPULA FAMILY.

Schedule pair (hours)		Parameter (and log likelihood)										
I_U	I_V	N_{obs}	Gaussian		Clayton		Gumbel		Frank		Student-t	
1	1	486										
1	4											
1	6	5										
1	8											
1	12	3000	0.183	(50.19)	0.140	(18.33)	1.150	(**73.72**)	1.217	(59.24)	14.02	(56.32)
4	4											
4	6											
4	8											
4	12											
6	6											
6	8											
6	12	30										
8	8											
8	12											
12	12	3000	0.296	(135.8)	0.438	(134.1)	1.269	(161.7)	2.444	(206.9)	3.255	(**242.1**)

TABLE II

EMPIRICAL COPULA FITTING FOR SCHEDULE PAIRS IN W9.

Schedule pair (hours)		Parameter (and log likelihood)										
I_U	I_V	N_{obs}	Gaussian		Clayton		Gumbel		Frank		Student-t	
1	1	1319										
1	4	125										
1	6	3000	0.220	(72.41)	0.130	(14.36)	1.197	(**128.4**)	1.373	(73.97)	6.622	(100.5)
1	8	3000	0.188	(52.55)	0.120	(12.86)	1.148	(**79.77**)	1.126	(50.85)	10.25	(63.72)
1	12	3000	0.170	(42.96)	0.078	(6.296)	1.164	(**95.99**)	1.004	(38.21)	4.871	(82.57)
4	4	15										
4	6	309										
4	8	563										
4	12	110										
6	6	3000	0.569	(582.6)	0.817	(327.0)	1.931	(902.1)	6.384	(**954.6**)	0.000	(19360)
6	8	3000	0.535	(502.0)	0.739	(276.4)	1.804	(756.9)	5.954	(**874.5**)	0.000	(39510)
6	12	274										
8	8	3000	0.552	(539.9)	0.865	(364.4)	1.828	(762.2)	6.506	(**989.1**)	2.596	(915.3)
8	12	192										
12	12	3000	0.561	(561.8)	0.916	(392.4)	1.817	(753.0)	6.519	(**1006**)	0.000	(40730)

policy, and that an appropriately balanced staffing resource is crucial in managing individual patient schedules effectively. These are novel findings in the field of healthcare informatics, and this new level of information has potential to help ward managers design staffing levels tailored to the individual requirements of a ward. Although our sample of study wards is limited to 20 in south Wales we can expect similar outcomes from other hospital wards across the UK that also use the same patient escalation frameworks.

Although there are many uncontrollable factors that encourage vital signs to be observed dependently in practice, a ward manager looking to improve patient care may consider reviewing specific operational dependencies that could potentially be detrimental to the delivery of patient care. For example, when 1 hour observations are not taken independently; observations on this schedule are often a minority and should never have a TTNO that straddles to the subsequent ward round, so forming strong copulas with long intervals, or even forming a statistically significant copula indicates room for improvements in staff resources. There is similar treatment with other less common intervals like 4 hour and medium-length intervals in ward group 1, where the rarity of these observations may be enough of a factor to encourage batching; perhaps clinical staff undertake other tasks between routine schedules and are caught under resourced for when they occur. Ward managers could find utility in evaluating changes in interval dependencies as part of broader retrospective reviews of policy adjustments, training programs implementation, or ward response to serious long-term events, like Covid-19, staffing shortages, or seasonal changes [11], [14].

Health boards may also be interested in the patient management efforts occurring across different sites. Currently it is impossible to make a fair comparison between two wards because of the vast number of practical variables in play, but it is often necessary for designing financial or resourcing strategies. Additional information that can characterise ward operations can at least provide operational context when evaluating ward performance. For example, Table III shows

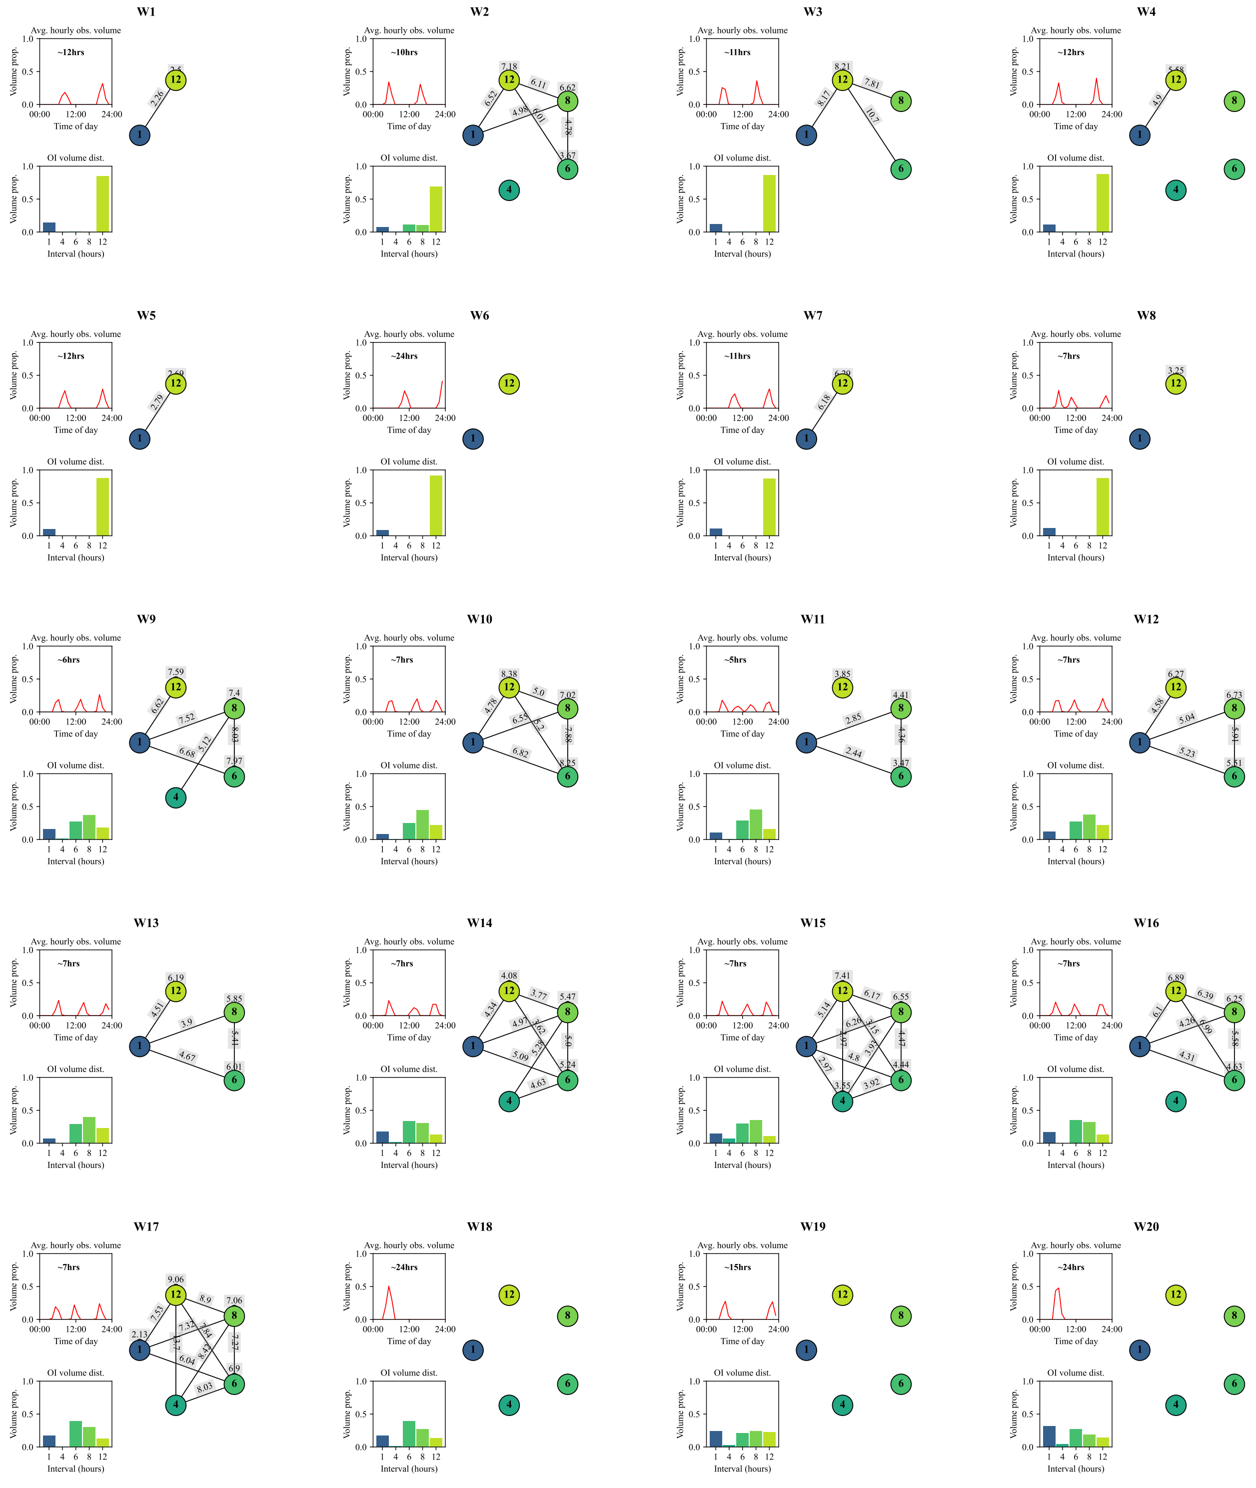

Fig. 6. θ parameters for the Frank copula fits derived from pairwise TTNO joint distributions for observations on different specified schedules. Edges are present for suitable copulas fits (where log-likelihood is greater than 50), and labelled with the exact value for θ. Nodes represent hazard distributions for each observation interval present in the ward. Upper inset plots show daily observation volume patterns for each ward (also shown in Figure 1), and lower inset plots illustrate volume distributions of observation intervals.

that wards W2, W8, and W11 share specialisms and are of similar size, and would suggest that W2 has the best timeliness conformance, however, in Figure 6 we see each ward using a different daily ward round schedule and operate with different interval dependency structures. Each of these wards would need independent resourcing strategies. On the other hand,

TABLE III
SPECIALISMS AND VITAL SIGN OBSERVATION VOLUMES FOR THE 20 STUDY WARDS.

Ward	Speciality	N_{obs}	on time
W1	Surgical	25629	63.06%
W2	Medical	28642	54.73%
W3	Cardiology	25476	68.59%
W4	Rehabilitation	12073	65.66%
W5	Cardiology	28466	65.51%
W6	Trauma & Orthopaedics	9335	75.48%
W7	Trauma & Orthopaedics	25137	67.37%
W8	Medical	25998	17.32%
W9	Gastroenterology	35359	42.22%
W10	Respiratory	37875	54.91%
W11	Medical	30659	30.67%
W12	Rehabilitation	39366	43.05%
W13	Care of the elderly	31461	50.56%
W14	Rehabilitation	44768	39.84%
W15	Rehabilitation	47671	40.15%
W16	Medical	45107	38.2%
W17	Rehabilitation	42605	38.05%
W18	Medical	54449	42.52%
W19	Medical	22249	23.45%
W20	Surgical	45716	31.2%

there are groups of very similar wards (rehabilitation wards W14, W15, W17, and cardiology wards W3, W5), for which the alignment in these additional metrics could give confidence to a health board that they could act as a references to one another when designing new resource strategies, policies, or training programs.

Alternatively, further studies may find an appropriately fitting copula family its expected value ranges for its θ parameter across all vital sign observation schedule parings essential information for synthetic vital sign observation data set creation. Synthetic data sets have wide use cases; practically this could be trialling potential operational adjustments, such as scheduling more daily ward rounds or a wider range of observation intervals, without having to place patients at risk. Otherwise, simulated data could find its way to supporting new dashboard information models that can aid with clinical staff or bed resource allocation methods.

V. Conclusion

The copula approach has proved to be an appropriate method to define and quantify the non-linear relationships in the TTNOs of any vital sign observation schedules pair at the point of the subsequent ward round. In 17 of the 20 study wards there was one or more significant dependency between standard observation intervals that reflect patterns in daily ward rounds and observation interval type distributions. It is now clear that modelling the TTNO of any given vital sign observation individually is ineffective for capturing the practical nature of observations management in hospital wards, and that understanding the current level of pairwise observation schedule dependency is essential knowledge for hospital ward stakeholders when design staffing resource strategies. This study is part of a path to smarter patient care by providing enhanced information sources that can be used as a basis of care policies tailored for the individual needs of any ward.

VI. Limitations and further work

Understanding how probable dependency between any two observations with specified scheduling is presents an opportunity to incorporate individual patient covariates, such as NEWS, Sepsis[6] labelling, and staff concern, to see how these may impact the shape of the TTNO marginals. There is potential in defining a complete model to forecast clinical staff vital sign observation management which incorporates all current patients within a ward and the interactions between those who have been observed within the same ward round. Following work will also incorporate multivariate data, specifically, when groups of observations are taken at a similar time (i.e., during ward rounds). Whether dependency correlates to the number of patient within a group, the intervals present, or by ward round heuristics is still unknown. This avenue could help hospital managers plan staff resources to address recurrent ward conditions that are most likely to encourage late or missed observations. We know the copula theory discussed in this paper does not reflect a complete solution to the dependence structures within vital sign observations management. Rather, this study serves as a demonstration to the level of frailty that exists within this data, and shows that it is an essential feature to include within any TTNO forecasting model.

Ethical Approval

The data request was approved by the Aneurin Bevan University Health Board research risk panel and received a favourable opinion from the Cardiff University School of Computer Science and Informatics research ethics committee.

Declaration of competing interest

All authors declare they have no competing interests.

Acknowledgements

The authors would like to acknowledge the assistance and cooperation of the Aneurin Bevan University Health Board for their provision of the supporting dataset and their informal commentary relating our findings to real ward operation.

[6]Sepsis occurs when the body's response to infection damages its own tissues and organs, where the chemicals released into the bloodstream to fight an infection trigger damaging inflammatory responses throughout the body. It is a medical emergency that requires immediate intervention, but early onset symptoms can be spotted (e.g., fever, elevated heart rate, and rapid breathing) and treatment administered from timely vital signs monitoring [20].

Author Contributions

R.I-S. processed the experimental data, performed computations, designed figures, and drafted the manuscript. S.A., B.N., and L.D.T. supervised the study, verified the analytical methods, and extensively assisted in interpreting the findings and preparing the manuscript.

Data Availability

The dataset used in this study was provided by Aneurin Bevan University Health Board and not available for public release by the authors due to its sensitive context. Aneurin Bevan University Health Board should be contacted for access.

References

[1] ABUHB. Aneurin bevan university health board deteriorating patient policy., 2017.

[2] Shiyu Zhou Akash Deep and Dharmaraj Veeramani. Copula-based multi-event modeling and prediction using fleet service records. *IISE Transactions*, 53(9):1023–1036, 2021.

[3] Ban Al-Sahab, Alan Leviton, Tobias Loddenkemper, Nigel Paneth, and Bo Zhang. Biases in electronic health records data for generating real-world evidence: An overview. *Journal of Healthcare Informatics Research*, 8(1):121–139, 2024.

[4] Wayne Nelson and. Hazard plotting for incomplete failure data. *Journal of Quality Technology*, 1(1):27–52, 1969.

[5] Jim Briggs, Ina Kostakis, Paul Meredith, Chiara Dall'ora, Julie Darbyshire, Stephen Gerry, Peter Griffiths, Jo Hope, Jeremy Jones, Caroline Kovacs, Rob Lawrence, David Prytherch, Peter Watkinson, and Oliver Redfern. Safer and more efficient vital signs monitoring protocols to identify the deteriorating patients in the general hospital ward: an observational study. *Health Soc Care Deliv Res*, 12(6):1–143, Mar 2024.

[6] Valérie Chavez-Demoulin and Paul Embrechts. Revisiting the edge, ten years on. *Communications in Statistics - Theory and Methods*, 39(8-9):1674–1688, 2010.

[7] Kelly M. Derby, Natalie A. Hartung, Sherry L. Wolf, Sherry Wolf, Heather L. Zak, and Laura K. Evenson. Clinical nurse specialist-driven practice change: Standardizing vital sign monitoring. *Clinical Nurse Specialist*, 2017.

[8] Samorodnitsky G. Embrechts P. and, Resnick S. and. *Living on the Edge*. RISK, 1998.

[9] Maurice J Frank. On the simultaneous associativity of f x y and x y f x y. *Aequationes Mathematicae*, 1978.

[10] Simon Gill, Bruce Stephen, and Stuart Galloway. Wind turbine condition assessment through power curve copula modeling. *IEEE Transactions on Sustainable Energy*, 3(1):94–101, 2012.

[11] P. Griffiths, K. Ball, J.and Bloor, D. Böhning, J. Briggs, C. Dall'Ora, A. D. Iongh, J. Jones, C. Kovacs, A. Maruotti, P. Meredith, D. Prytherch, A. R. Saucedo, O. Redfern, P. Schmidt, N. Sinden, and G. Smith. Nurse staffing levels, missed vital signs and mortality in hospitals: retrospective longitudinal observational study. *NIHR Journals Library*, 2018.

[12] C. Hands, E. Reid, and P. et al Meredith. Patterns in the recording of vital signs and early warning scores: compliance with a clinical escalation protocol. *BMJ*, 2013.

[13] Edward L. Kaplan, Edward L. Kaplan, E. L. Kaplan, Paul Meier, and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 1958.

[14] I. Kostakis, G. B. Smith, D. Prytherch, P. Meredith, C. Price, A. Chauhan, and Covid Portsmouth Academic ConsortIum For Investigating. Impact of the coronavirus pandemic on the patterns of vital signs recording and staff compliance with expected monitoring schedules on general wards. *Resuscitation*, 158:30–38, 2021.

[15] David Xianglin Li. On default correlation: A copula function approach. *SSRN*, 1999.

[16] F. McGain, M. A. Cretikos, D. Jones, S. Van Dyk, M. D. Buist, H. Opdam, V. Pellegrino, M. S. Robertson, and R. Bellomo. Documentation of clinical review and vital signs after major surgery. *Med J Aust*, 189(7):380–3, 2008.

[17] H. McGloin, A. Sk, and M. Singer. Unexpected deaths and referrals to intensive care of patients on general wards. are some cases potentially avoidable? *Journal of The Royal College of Physicians of London*, 1999.

[18] Sogol Moradian, Agnieszka I. Olbert, Salem Gharbia, and Gregorio Iglesias. Copula-based projections of wind power: Ireland as a case study. *Renewable and Sustainable Energy Reviews*, 175:113147, 2023.

[19] R Nelsen. *An Introduction to Copulas*. Springer, New York, 1999.

[20] NICE. Sepsis: recognition, diagnosis and early management., 2023.

[21] B. Noë, A. Bullock, J. Frankish, and L. D. Turner. Temporal patterns in vital sign recording within and across general hospital wards. *Resusc Plus*, 10:100247, 2022.

[22] David Oakes and David Oakes. Bivariate survival models induced by frailties. *Journal of the American Statistical Association*, 1989.

[23] G. N. Oliveira, L. S. Nogueira, and Dalmd Cruz. Effect of the national early warning score on monitoring the vital signs of patients in the emergency room. *Rev Esc Enferm USP*, 56(spe):e20210445, 2022.

[24] O.C. Redfern, P. Griffiths, A. Maruotti, A. Maruotti, A. Recio-Saucedo, A.R. Saucedo, and G.B. Smith. The association between nurse staffing levels and the timeliness of vital signs monitoring: a retrospective observational study in the uk. *BMJ Open*, 2019.

[25] F. Shewa, S. Endale, G. Nugussu, J. Abdisa, K. Zerihun, and A. Banbeta. Time to kidneys failure modeling in the patients at adama hospital medical college: Application of copula model. *J Res Health Sci*, 22(2):e00549, 2022.

[26] M Sklar. Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, 8:229–231, 1959.

[27] Yuankun Wang, Huiqun Ma, Dong Wang, Guizuo Wang, Jichun Wu, Jinyu Bian, and Jiufu Liu. A new method for wind speed forecasting based on copula theory. *Environmental Research*, 160:365–371, 2018.

[28] Y. Wei, M. Wojtys, L. Sorrell, and P. Rowe. Bivariate copula regression models for semi-competing risks. *Stat Methods Med Res*, 32(10):1902–1918, 2023.