neteera

August 4, 2025

Neteera System Efficacy Study in Predicting Hospitalizations or Adverse Events

Shai Shecter, PhD[1]; Alex Kazakov[1], PhD; Yair Deitcher, PhD[1]; Shimon Steinberger[1]; Shayna Hard[2], MSN, APRN, CNP; Hans Jenkins[2], DO, MBA; Mordy Eisenberg[2]; Isaac Litman[1]

1 Neteera
2 TapestryHealth

TABLE OF CONTENTS

Contributing Authors

Dr. Shai Shechter was responsible for data processing, data analysis, and statistical evaluation.

Dr. Alex Kazakov wrote the manuscript and prepared the figures, and contributed to data analysis and statistics.

Dr. Yair Deitcher and Shimon Steinberger jointly managed the research and development of the project, providing strategic oversight.

Shimon Steinberger and Dr. Yair Deitcher and jointly managed the research and development of the project, providing strategic oversight.

Shayna Hard, MSN, APRN, CNP, Director of Remote Patient Monitoring at TapestryHealth, served as the primary data representative from Tapestry, manually curating, aggregating, and delivering the relevant medical records for the study.

Dr. Hans Jenkins, DO, MBA, Chief Medical Officer at TapestryHealth, supervised and consulted on Shayna's work throughout the data contribution process.

Mordy Eisenberg, Co-Founder/Chief Growth & Product Officer at TapestryHealth, assisted in establishing the collaboration and operational coordination, working alongside Isaac Litman.

Isaac Litman, CEO of Neteera Technologies Ltd., provided organizational leadership and overall project support.

All authors reviewed and approved the final version of this report.

Abstract

Background

Early detection of health deteriorations in residents of healthcare facilities, such as nursing homes, can enable timely intervention and potentially prevent hospitalizations. This study evaluates the performance of the Neteera System (system used was the Neteera 130H-Plus), a contactless remote patient monitoring and clinical alert solution, in predicting hospitalization events among residents in TapestryHealth nursing facilities.

Methods

A retrospective analysis was conducted on 96 hospitalization events from an initial cohort of 612 monitored patients (across five facilities) over a four-month period. Continuous heart rate (HR) and respiratory rate (RR) data were collected by the Neteera System. Alerts triggered when vital signs deviated from an established baseline were computed for each patient throughout the entire monitoring period. Sensitivity (the fraction of hospitalizations preceded by an alert) and specificity (the fraction of non-event periods without false alerts) were computed for various alert parameter settings.

Results

The Neteera System demonstrated the ability to anticipate a significant portion of hospital admissions. Two alert configurations were analyzed, one optimized for sensitivity and the other one for specificity. In the sensitivity-oriented alert configuration, approximately 77% of all hospitalization events were preceded by system notifications, with a specificity of 71% (indicating modest false alarm rates). Sensitivity increased to 89% for a subset of 18 'definite' hospitalization events that occurred due to respiratory- or heart-related clinical causes. A specificity-optimized setting achieved 86% true-negative rate while still detecting about 55% of events (70% of 'definite' events). The trade-off between sensitivity and specificity is characterized, and two example settings are highlighted and analyzed. Notably, the first alert was issued 4.99 days (on average, with SD=2.15) before the hospital admission for the sensitivity-oriented setting, and 5.23 days (SD=2.16) for the specificity-oriented one.

Conclusions

Neteera 130H-Plus device shows promise in early identification of clinical deteriorations leading to hospitalizations. These results provide a baseline for optimizing alert algorithms. The findings underscore key trade-offs: more aggressive alerting can warn of more impending hospitalizations but increases false alerts, whereas stricter criteria reduce false positives but may miss some events. Ongoing and future work (Phase 2) will refine the alert parameters to improve this balance, and Phase 3 will introduce multivariate "critical" alerts. Ultimately, the system aims to reduce adverse events and emergency hospitalizations by enabling proactive care interventions.

Introduction

Unplanned hospitalizations of residents in nursing facilities are common and often preventable. A study of 11 nursing homes over five years revealed that reducing avoidable hospitalizations resulted in over $2.6 million in recaptured revenue for the facilities and more than $31 million in savings for payers.[1] These hospital transfers are not only costly but dangerous - readmitted patients face a substantially higher mortality risk (four-fold increase at six months post-discharge).[2]

Early identification of patient deterioration is critical in reducing avoidable hospitalizations.[3,4] Changes in vital signs are often the first visible signs of decline; an abnormal respiratory rate (RR) is a well-recognized early indicator of infection and other acute issues.[3,5,6] However, in low-acuity settings like nursing homes or non-ICU units in hospitals, vital sign monitoring is typically intermittent and infrequent (e.g. periodic nurse rounds). These facilities usually measure respiratory rate by manual counting, an error-prone and infrequent process. [6] This gap in monitoring can lead to delays in recognizing deterioration. Studies have demonstrated a direct link between delayed transfer to a higher level of care, such as an intensive care unit (ICU), and increased hospital mortality.

Advances in remote patient monitoring technology offer a solution. The Neteera System includes a contactless sensor device that continuously tracks patients' cardiopulmonary signals without wires or wearables *(Figure 1a)*. This wall-mounted device detects micro-movements of the patient's skin due to heartbeats and breathing, deriving continuous heart rate and respiratory rate readings in real time *(Figure 1b)*. Importantly, the system analyzes these vital sign trends relative to each patient's baseline and generates alerts when deviations exceed clinically defined thresholds (Figure 2). This continuous monitoring paradigm enables 24/7 vigilance, closing the several-hour gaps (especially at night) when residents are typically unmonitored by staff.

Tapestry Health, a provider of telemedicine and remote care services, has partnered with Neteera to implement this technology in multiple long-term care centers. The Neteera System provides clinical teams with early warnings, enabling them to intervene on-site before a condition worsens to the point of requiring hospitalization. Specifically, the goal of Phase 1 (reported here) is to quantify the predictive power of the current generation alert algorithm in identifying upcoming hospitalizations, using retrospective data from an initial cohort of residents. A hospitalization event is used as a proxy for a severe adverse health event, under the premise that many urgent hospital transfers are preceded by physiological deterioration that the system could potentially detect early. [7,8]

This study was conducted over an initial cohort of 612 patients. The analysis of performance over the initial cohort is divided into three phases, as follows: Phase 1 reports system performance using the current alert parameters retrospectively; Phase 2 will involve fine-tuning these alert parameters based on Phase 1 findings;

and Phase 3 will introduce novel, multi-parameter “critical” alerts that combine several vital sign factors for improved overall performance. Here we present the Phase 1 results for the initial cohort, including background on the data collected, the alerting mechanism, analysis methods, and performance outcomes (sensitivity and specificity trade-offs). The implications of these findings are discussed, the practical trade-offs for deployment are considered, and the planned next steps to enhance the system’s predictive capabilities are outlined.

Figure 1

The Neteera 130H-Plus device monitors vital signs remotely and continuously.

Figure 1a

Neteera 130H-Plus is installed on a wall mount up to 2.5 meters (8 ft.) from the patient’s bed.

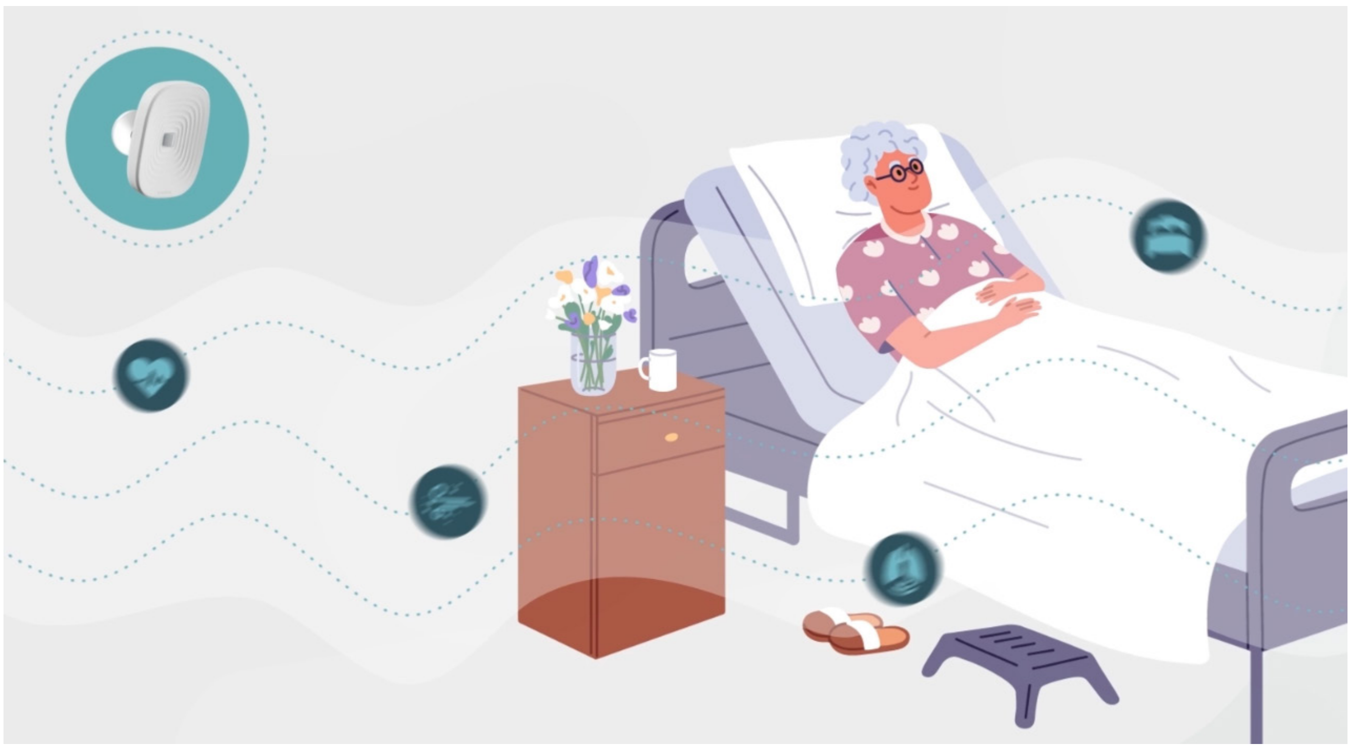

Figure 1b

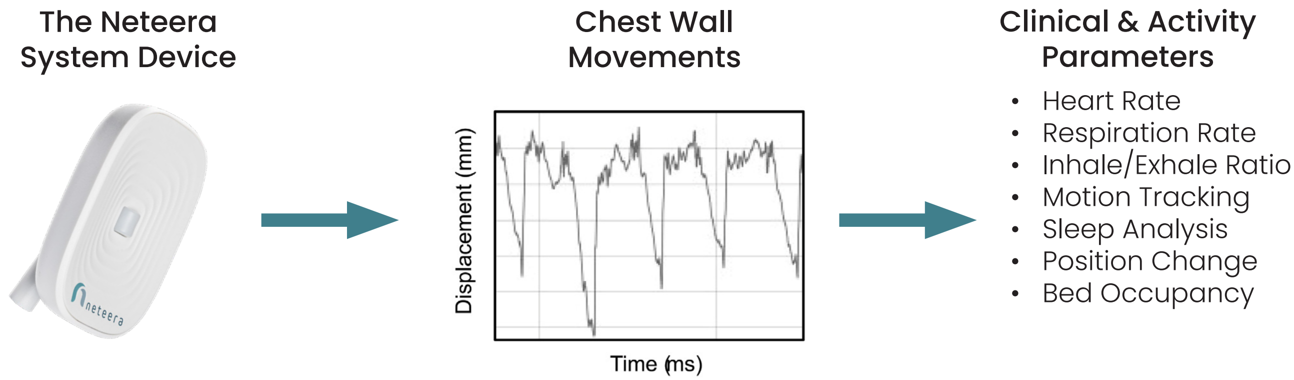

a: *The Neteera 130H-Plus is installed on the wall up to 2.5 m (8 ft.) from the resident's bed, remotely and continuously tracking the patient's micro-motions through clothing and bedding.*

b. *The sensor's raw displacement signal (example trace, center) is converted in real time into cardiopulmonary metrics (heart rate, respiratory-rate, inhale/exhale ratio, respiratory amplitude), sleep-related parameters (sleep duration and stages, position change, motion tracking, bed occupancy, bed exit, and pre-screening for obstructive sleep apnea which is EU CE-certified by NB 2797 with FDA review pending) and others. These derived vital and activity sign streams form the basis of the personalized alert algorithm evaluated in this study.*

Medical Data Collection

Study Cohort and Setting

Data were collected from five nursing facilities of TapestryHealth's remote monitoring program in the US. These facilities installed Neteera's contactless vital sign devices for continuous patient monitoring. The study period spanned from November 1, 2024, to February 28, 2025. During this four-month window, 612 residents were continuously monitored by the Neteera System and thus eligible for this analysis.

Patient and Event Counts

Among the 612 monitored patients (the initial cohort), 75 individuals experienced one or more hospitalizations during the study period. In total, 96 distinct hospitalization events were recorded in this cohort. The hospitalizations ranged in cause and urgency, and each was documented through facility records (e.g. emergency department transfers or inpatient admissions).

Recording Setup and Alerting System

Contactless Monitoring of Vital Signs

All facilities in the study were equipped with Neteera's 130H-Plus contactless vital signs monitoring devices (Figure 1a). Each device is mounted in the patient's room (typically on a wall or ceiling) and detects minute body surface motions from cardiac pulses and respiration (Figure 1b). This technology allows **ambient, continuous, non-intrusive monitoring** of heart rate and respiratory rate through clothing and bedding, without wires or wearables on the patient. The system automatically filters and processes the micro-motion signals to output real-time HR and RR measurements for each patient, providing a constant feed of vital sign data.

Alert Algorithms (Threshold vs. Baseline)

The alert system in the Phase 1 study was based on a personalized "drift from baseline" approach that tracked temporal trends of each vital sign. An alert was issued if the ratio of these averages exceeded a predefined bound. The alert system continuously computes long- and short-term averages of the vital sign, and issues an alert if the ratio of averages exceeds a predefined boundary. At every time step, the device computes:

- **Baseline vital sign average:** the running mean of the vital sign over the preceding three days (intended to represent the patient's habitual resting level).
- **Current vital sign average:** the mean over the most recent four hours.

An alert is generated when the current window mean exceeds the baseline mean by more than a configurable percentage (e.g., 25%). Baseline alerts are tailored to individual trends in vital signs, accounting for each patient's typical resting state. Additionally, these alerts maintain resilience against temporary fluctuations, as current averages are determined based on several hours of collected data.

In contrast to Neteera's baseline alert, a fixed-threshold alert triggers whenever a spot vital sign measurement exceeds a pre-defined absolute value. For example, a heart rate (HR) alarm may be set at 110 bpm; any reading that exceeds 110 bpm immediately issues an alert.

Examples of threshold and baseline alerts are shown in Figure 2, where heart-rate readings are provided for two patients. Each dot corresponds to an average vital sign reading reported during a five-minute period. The red vertical line marks the onset of the alert. A grey rectangle marks the hospitalization event. Figure 2a illustrates an alert issued by the 'threshold' logic. A heart-rate alert is triggered a day before hospitalization when the patient's heart-rate readings exceed 110 beats per minute.

Figure 2b shows multiple alerts triggered by the 'baseline' logic. The adaptive rule triggers as soon as the ratio of the long-term to the short-term means exceeds 25%. In practice, the alarms were issued at 81.8 bpm, 84.0 bpm, and 90.1 bpm during the fourth day, third day, and one day before hospitalization. The bpm rates were well below the naive 110 bpm threshold, yet the accelerating heart rate was captured, culminating in a hospitalization.

Parameter exploration in the present study: In Phase 1, a grid search was performed over the alert parameters (baseline ratio percentage, vital sign channel) to quantify how sensitivity, specificity, and lead time were modulated. The configuration illustrated in Figure 2 (three-day baseline, four-hour current window, ratio = 25 %; monitored vital sign = RR) is one representative point in that space.

Figure 2

Personalized vital-sign deviations trigger early alerts days before an eventual hospitalization.

a. *Time series data for a single resident is shown as five-minute-averaged dots during the last five days prior to hospitalization. The resident's long-term heart rate baseline is stable until the last day prior to the hospitalization event, where the heart rate crossed the predefined threshold of 110 bpm, issuing an alert (red vertical line). The vertical gray rectangle after the mark '0' indicates the resident's transfer to the hospital.*

b. *Baseline-alert logic is applied to another patient's heart-rate readings, covering seven days prior to hospitalization. Every alert event, marked by a red line, was issued since the "current" average of heart rate over the last four hours exceeded the defined threshold of 25% over the average of the last three days.*

Methods

Study Design

This is a retrospective observational study evaluating the predictive alert performance of the Neteera System. Historical vital-sign alert data was paired with recorded hospitalization events in the initial cohort. Because the study is retrospective, no changes to resident care were made based on the alerts during the period. Analysis is based on what the system would have detected relative to what occurred.

Sensitivity and Specificity Definitions

Sensitivity is defined as the proportion of hospitalization events for which at least one alert was raised in the days preceding the event. Specifically, an alert was considered successful in "predicting" a hospitalization if it occurred within a window of up to seven days before hospital admission. Any hospital event with no alert in the prior seven days was counted as a missed event (false negative). This relatively long prediction window (seven days) was chosen to capture cases where vital signs begin trending abnormally days in advance of acute decompensation.

For **specificity**, the frequency of alerts that occurred without a subsequent hospitalization event is assessed. Using a daily-resolution approach, a false-positive was defined as an alert generated for a single day's records from a patient who was not hospitalized. Specificity was then calculated as the fraction of no-alert days out of all single-day recordings from the non-hospitalized patient group. Thus, high specificity means that there are very few alerts on days with no ensuing hospitalization events. Some alerts might reflect genuine issues that are resolved in the facility without hospitalization, but for this hospital-focused analysis, they are counted as false positives if no admission followed.

Using these definitions, sensitivity and specificity for the entire cohort were computed. Calculations were performed for two categories of events: **(1)"All hospitalizations,"** and **(2) "Definite" hospitalizations**—the latter referring to a subset of hospitalization events identified as likely due to respiratory- or heart-related clinical deterioration (such as cardiopulmonary events, shortness of breath, etc.), as opposed to pre-scheduled or non-medical admissions. The "definite" subset (18 events of the 96 total) represents those cases where an early warning from heart or respiratory vitals should theoretically be possible.

Parameter Search

The effect of **varying alert ratio parameters** on sensitivity and specificity was explored. A parameter search was conducted over selected ratio values. In practice, this meant re-running the alert detection logic on the recorded vital sign data under different ratio criteria (e.g., requiring a 25% rise from baseline vs 30% vs 35%, etc.) and recalculating sensitivity/specificity for each scenario. We varied three parameters over a small set of values, while keeping other factors (such as the baseline window) fixed, to map out the performance trade-off curve. See Table 1 for the list of parameters and the sweep ranges in our parameter search.

Table 1

A list of variables that were changed during the parameter search. Most variables of the alert policy were kept constant and are not listed here.

VARIABLE NAME	TESTED VALUES	AFFECT ON ALERT LOGIC
Ratio (Current Baseline)	20% 25% 30% 35%	Defines the deviation in % of vital sign readings from the baseline in the last four hours over the previous three days.
Monitored Vital Sign	Respiration Rate Heart Rate	Each vital signal has its own dynamics, resulting in different performances. Combination of both the HR and RR alerts increases the alert rate, typically leading to increased sensitivity at the expense of a decrease in specificity.
Alert Time Window Prior to Hospitalization	3 Days 7 Days	Time window (in days) before hospitalization during which an alert is considered a true positive event.

This systematic sweep allowed us to identify Pareto-optimal configurations - points at which sensitivity cannot be improved without sacrificing specificity, and vice versa. By analyzing this trade-off "manifold," two representative operating points were selected for detailed focus: one optimized for **high specificity**, and one for **high sensitivity**.

Analysis of Alert Timing

In addition to sensitivity and specificity, the timeliness of the alerts was analyzed. For every hospitalization that was successfully preceded by an alert (a true positive), the time difference in days between the first alert and the subsequent hospital admission was calculated. The mean and standard deviation of these "days before hospitalization" for each alert configuration were computed. Analysis was performed to quantify the average warning window provided by the Neteera System, as a longer lead time offers a greater opportunity for clinical intervention.

The study was conducted under appropriate data use agreements, and all patient data were handled in compliance with privacy regulations.

Results

Through the parameter search, the influence of adjusting the alert ratio on performance was quantified. Figure 3 plots the sensitivity versus specificity for each tested configuration, forming a characteristic trade-off curve. As expected, there is an inverse relationship: making the alert trigger more easily (lowering the ratio for deviation) improves sensitivity (more events are caught, the points appear further to the right) but reduces specificity (more false alarms, the points drift to the bottom), while a stricter trigger (higher ratio) does the opposite.

Figure 3

Trade-off curve from the parameter search of alert parameters.

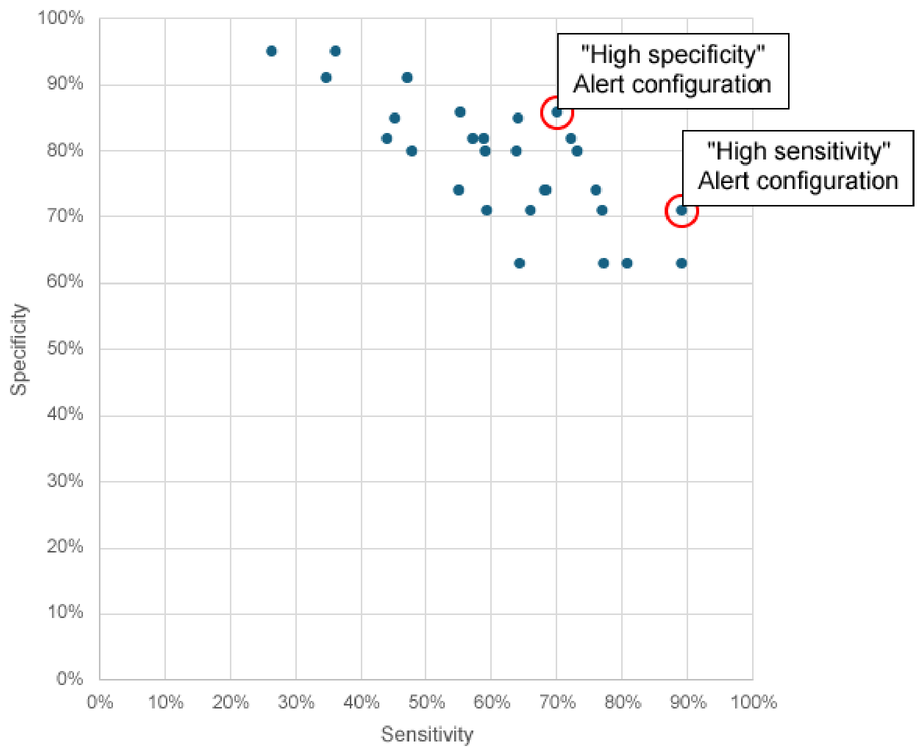

Each blue dot represents one of the tested alert configurations (different baseline deviation ratios, vital sign sources, and prediction windows), plotted by its resulting sensitivity (ability to detect a hospitalization) on the x-axis and specificity (fraction of quiet days with no false alert) on the y-axis. The cluster forms a downward-sloping Pareto frontier. loosening the trigger criteria improves sensitivity but inevitably lowers specificity. Two configurations of practical interest are highlighted with red circles: a "high specificity" setting (55% sensitivity / 86% specificity) and a "high sensitivity" setting (77% sensitivity / 71% specificity). These points illustrate the operational choices available to facilities depending on whether minimizing missed events or minimizing false alarms is the priority.

Analysis also revealed that when the system issued a pre-hospitalization alert, it did so with a substantial lead time, which we measured in days before the hospitalization event.

Two notable configurations on this curve were identified:

- **High Sensitivity Configuration:**
  Ratio = 25% above baseline, only respiration alerts

  Out of the 96 total hospitalization events in the monitored cohort, the system generated at least one alert in the seven days prior to 73 of those events, yielding an overall sensitivity of 77%. This configuration is shown by the red circle on the right in Figure 3. Performance metrics are illustrated in Figure 4a. Note that only respiration events were considered as an alert. When focusing on the 18 "definite" deterioration-related events, 16 were preceded by alerts, corresponding to a higher sensitivity of 89%. The left two bars in Figure 4a. This indicates that most of the truly acute medical events were picked up by the vital sign alerts. Specificity was 71%, as shown by the blue bars in Figure 4a.

  For "All hospitalizations" events, on average, the **first alert appeared 4.99 days** (SD=2.15) before the hospital admission. For the most critical "definite" events, this lead time was even longer, averaging **5.6 days** (SD=1.92). Note that the "definite" events were alerted an average of 15.8 hours earlier than the combined group.

- **High Specificity Configuration:**
  Ratio = 35% above baseline, only respiration alerts

  Specificity reached 86% (only 14% of non-event periods generated alerts) with a sensitivity of 70% for the "definite" hospitalizations and 55% for all hospitalizations. This configuration is indicated by a red circle in Figure 3 and illustrated in Figure 4b. As in the previous setting, only respiration rate alerts were considered.

  For "all hospitalizations" events, the f**irst alert appeared 5.23 days** (SD=2.16) before the hospital admission on average. For the most critical "definite" events, this lead time was even longer, averaging **5.82 days** (SD=1.82). Note that the "definite" events were alerted on average 15.8 hours earlier than the combined group.

These outcomes are summarized in Table 2 on the following page, which lists the performance metrics for the three highlighted configurations. The table shows the sensitivity and specificity for both "**all events"** and "**definite events"** subsets under each configuration.

Table 2

Performance of the Neteera System under two representative configurations. "Ratio" refers to the vital sign deviation from patient baseline required to trigger an alert. Sensitivity is the percentage of hospitalizations that were preceded by an alert (with counts in parentheses), evaluated for all events and for the subset of definite deterioration-related events. Specificity is the percentage of days without any alert from non-hospitalized patients.

ALERT CONFIGURATION	VITAL SIGN	RATIO DEVIATION FROM BASELINE	SENSITIVITY: ALL EVENTS	SENSITIVITY: "DEFINITE EVENTS"	SPECIFICITY: NON-EVENT DAYS
High Sensitivity (Respiration Alert)	Respiration	25%	77% (74/96 Events) **Avg. Warning Time (Days):** 4.99 Days (SD 2.15)	89% (16/18 Events) **Avg. Warning Time (Days):** 5.65 Days (SD 1.92)	71%
High Specificity (Respiration Alert)	Respiration	35%	55% (53/96 events) **Avg. Warning Time (Days):** 5.23 Days (SD 2.16)	70% (12/18 events) **Avg. Warning Time (Days):** 5.82 Days (SD 1.82)	86%

Figure 4

Sensitivity-specificity trade-offs at two representative alert settings.

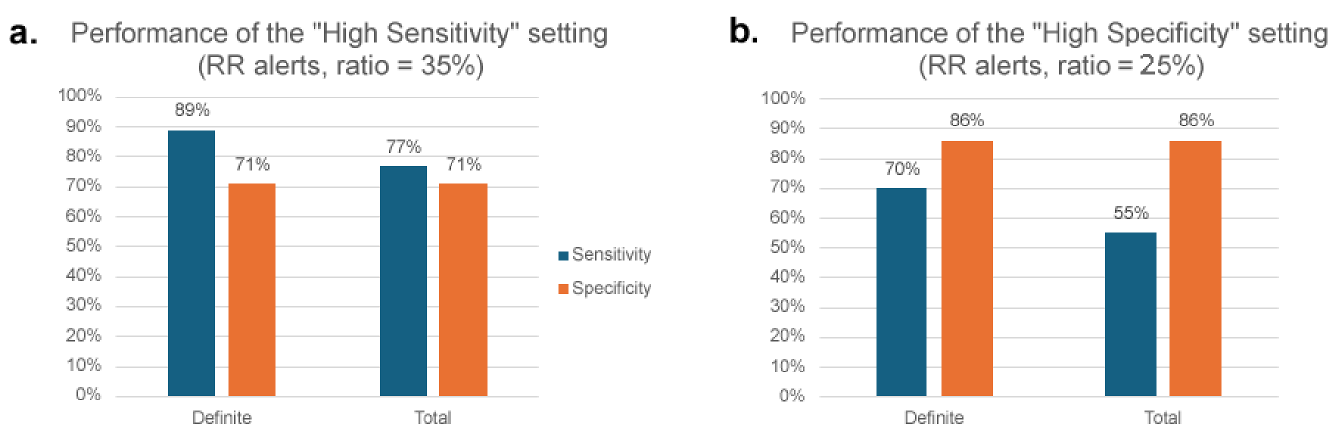

a. *High-sensitivity configuration: Respiration-rate alerts triggered when the four-hour average deviates ≥ 25% beyond the baseline. This captures 89% of "definite" hospitalizations and 77% of all hospitalizations, at the expense of a lower specificity of 71%.*

b. *High-specificity configuration: The threshold is tightened to a 35% deviation, raising specificity to 86% while sensitivity falls to 70% for definite events and 55% overall. Bars show the percentage of events (blue) or non-event days (orange) correctly identified, underscoring the practical balance between catching more deteriorations and minimizing false alerts.*

Discussion and Summary

Parameter Search Process

As illustrated in Figure 3, the resulting cluster of points revealed an inverse correlation, reflecting an inherent trade-off: lowering the ratio monotonically increased sensitivity and decreased specificity. No single configuration dominated all others; instead, different ratios yielded different compromises. This parameter search was a relatively coarse search (Table 1 lists the ranges of the examined variables), but it was sufficient to find near-optimal settings, optimized for sensitivity or specificity. Additional parameters (e.g., altering the baseline window duration or adding combined HR+RR criteria) will be explored in Phase 2. However, the current analysis already indicates that the existing algorithm, with an appropriately chosen ratio, can achieve high performance for single-sensor (HR, RR) alerting.

Interpretation of Findings

Phase 1 results demonstrate that the Neteera System, using a baseline-relative alert algorithm, can predict a substantial portion of impending hospitalizations in nursing facilities. In practical terms, the majority of serious health deteriorations leading to hospital transfers could be detected in advance by observing changes in heart or respiratory rates (Figure 4). This is a meaningful level of predictive power, considering the complexity of the various health events in nursing facilities. Not all hospitalizations are preceded by vital sign abnormalities. For example, some events, like sudden traumatic injuries or scheduled medical treatment, may not show a gradual trend, which inherently limits the maximum achievable sensitivity. The analysis of "definite" vs "all" events supports this: after excluding events less likely to be preventable by monitoring, the system's sensitivity rises (e.g. +12% from 77% to 89% for the 'High Sensitivity' setting, and +15% from 55% to 70% for the 'High Specificity' setting. See Table 2), suggesting that for the majority of predictable medical deteriorations, the system can flag the majority of them. Specificity of 71% and 86% (for the sensitivity- and specificity-optimized settings, correspondingly) indicates a moderate rate of false alerts, which is a reasonable trade-off in exchange for early warnings, but still implies that staff will need to respond to some alerts that do not result in major events.

Clinical Significance

An alert sensitivity of 55% or 77%, depending on the alarm system setting, means that a majority of hospitalizations had an advance warning sign from the system. From a clinical operations perspective, this could translate to opportunities for early intervention: if staff act on those alerts (e.g., assessing the patient, administering treatment, or calling a physician), some hospital transfers might be averted, or at least delayed. It is important to note that the study did not measure whether intervention on alerts would prevent admissions - that requires a prospective trial

- but evidence from similar deployments is encouraging. For instance, continuous monitoring of respiratory rate in care settings, such as skilled nursing facilities (SNFs) has shown the potential to initiate timely treatments that prevent hospitalizations for conditions like pneumonia.[3,9]

A key finding of this study is not just that alerts preceded hospitalizations, but when the alerts occur. An average warning window of over five days is a clinically significant period. This amount of lead time moves the intervention opportunity from the realm of emergency response to that of proactive care planning. This finding directly counters the concern that alerts might only happen immediately before an acute event, when interventions are less likely to be effective. Instead, it reinforces the system's value as an effective early warning tool. Furthermore, the data reveals that the system provides the earliest warnings for the most clinically relevant events. The average lead time for "definite" hospitalizations in the 'High Sensitivity' setting, for example, was 5.65 days, which is notably longer than the 4.99 days for events classified as "not relevant." This finding is crucial, suggesting that when primary cardio-respiratory issues drive deterioration, the very signals the system monitors, the physiological changes are detectable sooner and more reliably.

The **key trade-off** revealed by the results of this study is between missing events and raising false alarms. A higher sensitivity configuration (77–89% event detection) could maximize the chance of early intervention. Applied to patient safety, every additional warning could mean a life saved or a hospitalization avoided. However, that setting might contribute to alarm fatigue if too many alerts turn out to be benign. A study on the implementation of the National Early Warning Score (NEWS) found that alert systems were often ignored by front line staff and had no significant impact on clinical outcomes, highlighting that a predictive algorithm alone is insufficient without effective clinical integration and user acceptance.[10] On the other hand, a high specificity setting (86%) significantly reduces unnecessary alarms but would fail to alert in 22% of true deteriorations. The optimal point may also vary by facility or patient population: a high-acuity unit with ample staff might lean toward more sensitivity (accepting more alerts to ensure few events slip by), whereas a low-staff or lower-acuity environment might favor higher specificity to avoid overburdening the team with false positives.

Limitations

This analysis has several important limitations. The specificity calculation may be an overestimate, as some alerts counted as false positives could have reflected genuine health issues that were successfully managed by staff, thereby preventing hospitalization. Additionally, our cohort, while sizeable (612 patients), was still limited to certain facility types and geographies under the program, and only 96 hospitalization events.

Other vital signs or parameters were not incorporated into the study; only heart and respiratory rates were used. Other vitals and activity tracked by the Neteera System, including activity levels, personal alert history, medical records, and conditions, etc.) were not integrated into the alert logic at this phase of the study.

Overall, the findings confirm that continuous vital sign monitoring, combined with an intelligent alerting algorithm, can serve as an effective early warning system for adverse health events in care settings, including healtcare facilities such as SNFs. In this initial cohort, the Neteera 130H-Plus device achieved a balance of sensitivity and specificity comparable to that of early warning score systems in hospitals. For example, certain hospital early warning scores attain sensitivities around 60 - 70% for deterioration at false alarm rates of 20 - 30%, which is in line with what is seen here, albeit in a different setting.[7,11] The advantage of the Neteera System is its passive, continuous monitoring, which provides numerous data points to detect subtle trends that human monitoring might miss. There were clear cases where the respiratory rate increased days before clinical recognition of infection, precisely the kind of pattern this system is designed to detect. By quantitatively mapping performance, this study provides evidence to stakeholders (clinicians, administrators, and payers) that Neteera's technology is effective and continued development will enhance its effectiveness further.

Future Steps

Initial Cohort, Phase 2: Alert Parameter Refinement

Implementation of more thorough research of the parameters producing alerts from vital signs. The same specificity and sensitivity metrics will be used to evaluate the performance of each specific setting. The best settings will be chosen for different scenarios: Optimized for sensitivity, for specificity, and a balanced mode.

Initial Cohort, Phase 3: Multi-Variate "Critical" Alert Development

In Phase 3, the system will be augmented with more advanced alert mechanisms to synthesize multiple signals, in other words, a "critical meta-alert". Extensive research supports the claim that predictive models using a rich array of data from electronic medical records (EMRs) outperform traditional, simpler early warning scores 7. Studies directly comparing standard prediction rules with machine learning-based scores have found that machine learning approaches can more accurately identify patients at high risk of readmission while flagging significantly fewer patients, thereby improving the efficiency of clinical interventions. [12]

Combining heart rate, respiratory rate, other vital signs and sources of information, and patient medical records are planned. Concurrent tachycardia with tachypnea, for example, might be a stronger indicator of deterioration than either individually. Neteera's roadmap includes expansion of vital sign capabilities, such as arrhythmia. Using these additional data streams, algorithms will be developed to detect complex signatures of patient decline (such as heart rate variability changes, breathing irregularity, or a cluster of moderate vital sign shifts that together indicate health deterioration). Statistical models will be employed on the historical dataset to identify patterns that preceded adverse events, informing the meta-alert criteria. The intended outcome of Phase 3 is an enhanced alert system capable of catching cases that alerts based on single vital sign ratios cannot.

Larger Cohort: Large-Scale Validation (Second Cohort)

Following refinements from the initial cohort of patients, deployment will scale up to a larger cohort of patients. The purpose of this follow-up study is to validate the system's efficacy generalized at scale and to formally measure outcomes. A broader study will be conducted to assess metrics such as: reduction in hospitalization rates compared to baseline or control groups, reduction in emergency room visits, impact on patient health outcomes (mortality, length of stay, complication rates), and economic benefits (cost savings from avoided hospitalizations). With a large sample, performance can be stratified by subpopulations. For example, does the alert system work equally well with dementia patients versus cardiac patients, or are there patient factors that necessitate different ratio tuning? Operational metrics, such as alert frequency per patient-day and staff response times, will be monitored to ensure the system remains practical as it scales. Importantly, the study over a larger cohort of patients will involve close collaboration with medical advisory boards and clinical stakeholders to iterate on any workflow changes needed to effectively integrate the alerts into care processes (for instance, refining the protocol for responding to an alert).

References

1. Rantz, M. et al. Financial and Work-flow Benefits of Reducing Avoidable Hospitalizations of Nursing Home Residents. The Journal of Nutrition, Health and Aging , 971–978 (2021).

2. Shaw, J. A. et al. Thirty-Day Hospital Readmissions: A Predictor of Higher All-Cause Mortality for Up to Two Years. Cureus (2020) doi:10.7759/cureus.9308.

3. Lauteslager, T. et al. Detecting early signs of deterioration and preventing hospitalizations in skilled nursing facilities using remote respiratory monitoring. Respiratory Medicine Case Reports 0, 102044 (2024).

4. Ouslander, J. G. et al. Interventions to Reduce Hospitalizations From Nursing Homes: Evaluation of the INTERACT II Collaborative Quality Improvement Project. J American Geriatrics Society 9, 745–753 (2011).

5. Cei, M., Bartolomei, C. & Mumoli, N. In-hospital mortality and morbidity of elderly medical patients can be predicted at admission by the Modified Early Warning Score: a prospective study. International Journal of Clinical Practice 63, 591–595 (2009).

6. Moore, L. J. et al. Validation of a Screening Tool for the Early Identification of Sepsis. Journal of Trauma: Injury, Infection & Critical Care 66, 1539–1547 (2009).

7. Kipnis, P. et al. Development and validation of an electronic medical record-based alert score for detection of inpatient deterioration outside the ICU. Journal of Biomedical Informatics 64, 10–19 (2016).

8. Escobar, G. J. et al. Early detection of impending physiologic deterioration among patients who are not in intensive care: Development of predictive models using data from an automated electronic medical record. Journal of Hospital Medicine 7, 388–395 (2012).

9. Burke, R. E., Xu, Y. & Rose, L. Skilled Nursing Facility Performance and Readmission Rates Under Value-Based Purchasing. JAMA Netw Open , e220721 (2022).

10. Bedoya, A. D. et al. Minimal Impact of Implemented Early Warning Score and Best Practice Alert for Patient Deterioration. Critical Care Medicine 47, 49–55 (2019).

11. Escobar, G. J. et al. Automated Identification of Adults at Risk for In-Hospital Clinical Deterioration. N Engl J Med 383, 1951–1960 (2020).

12. Morgan, D. J. et al. Assessment of Machine Learning vs Standard Prediction Rules for Predicting Hospital Readmissions. JAMA Netw Open , e190348 (2019).

neteera.com

Produced in cooperation with

TapestryHealth™

The Neteera device should be used only as specified in the Instructions for Use (IFU) and according to its intended use.