

What Clinicians Want: Contextualizing Explainable Machine Learning for Clinical End Use

Sana Tonekaboni^{*1,2,4}

STONEKABONI@CS.TORONTO.EDU

Shalmali Joshi^{*2}

SHALMALI@VECTORINSTITUTE.AI

Melissa D. McCradden^{2,3,4}

MELISSA.MCCRADDEN@SICKKIDS.CA

Anna Goldenberg^{1,2,4}

ANNA.GOLDENBERG@VECTORINSTITUTE.AI

¹ *Department of Computer Science*

University of Toronto, Toronto, Canada

² *Vector institute for artificial intelligence*

Toronto, Ontario, Canada

³ *Department of Bioethics*

The Hospital for Sick Children, Toronto, Canada

⁴ *Department of Genetics & Genome Biology*

The Hospital for Sick Children, Toronto, Canada

Abstract

Translating machine learning (ML) models effectively to clinical practice requires establishing clinicians' trust. Explainability, or the ability of an ML model to justify its outcomes and assist clinicians in rationalizing the model prediction, has been generally understood to be critical to establishing trust. However, the field suffers from the lack of concrete definitions for usable explanations in different settings. To identify specific aspects of explainability that may catalyze building trust in ML models, we surveyed clinicians from two distinct acute care specialties (Intensive Care Unit and Emergency Department). We use their feedback to characterize when explainability helps to improve clinicians' trust in ML models. We further identify the classes of explanations that clinicians identified as most relevant and crucial for effective translation to clinical practice. Finally, we discern concrete metrics for rigorous evaluation of clinical explainability methods. By integrating perceptions of explainability between clinicians and ML researchers we hope to facilitate the endorsement and broader adoption and sustained use of ML systems in healthcare.

1. Introduction

For clinical Machine Learning (ML), lack of model robustness (Papernot *et al.*, 2016), complexity of clinical modeling tasks (Ghassemi *et al.*, 2018b), and high stakes (Vayena *et al.*, 2018) are some of the technical barriers to practical adoption. Additionally, even

a highly accurate ML system is not necessarily sufficient in and of itself to be routinely utilized and endorsed by clinical staff (Bedoya *et al.*, 2019; Guidi *et al.*, 2015). The uptake and sustained use of ML in healthcare has been more challenging than anticipated, as observed through piloting models for cardiac arrest (Smith *et al.*, 2013) and sepsis (Masino *et al.*, 2019; Elish, 2018). This brings about the question of how truly reliable ML systems can earn clinicians’ trust for sustained use and positive clinical impact.

Currently, many rule-based assistive tools, such as different early warning scores (Prytherch *et al.*, 2010; Vincent *et al.*, 1996; Duncan *et al.*, 2006), have been deployed in clinical settings. By their very nature, rule-based algorithms are not considered opaque and in need to be endowed with auxiliary explanations of model outcomes. However, recently developed ML based tools that aim to improve on rule based systems are not considered to have the same level of accepted transparency. Therefore, understanding ML model behaviour beyond conventional performance metrics has become a necessary component of ML research, especially in healthcare. This has led to the development of “interpretable” or “explainable” machine learning models (Gunning, 2017) as a measure to overcome barriers of trust and adoption. Unfortunately, as it stands right now this is a rather ill-defined problem (Lipton, 2016). Existing work suffers from significant criticism due to the lack of objective definitions of what would validate a model in terms of explainability in the context of clinical practice, while only limited works focus on evaluating the quality and usability of the proposed explanation methods for the target end user (Doshi-Velez and Kim, 2017). The ML community has largely resorted to developing novel techniques for model explanations (Guidotti *et al.*, 2018) that are often insightful only to the algorithmic experts (Miller, 2018). Recent commentators have suggested that perhaps the ‘explainability’ literature as it stands requires a pragmatic lens to better suit the needs of the end users who are ultimately the ones who will decide if the model survives in the real world or not (P  ez, 2019).

For the purpose of this manuscript, we define “Explainability in ML for Healthcare” as a set of measurable, quantifiable, and transferable attributes associated with an ML system targeted for clinicians to calibrate model trust. Our goal is to synthesize a pragmatic notion of explainability for clinicians and effectively inform on-going research efforts in clinical ML at our institution. The integration of clinical ML in hospital settings is prompting the application of qualitative methodology in a novel context (Elish, 2018). While previous work has assessed clinician use and acceptance of ML tools retrospectively and through quantitative surveys (Guidi *et al.*, 2015), we are among the first to investigate prospectively users’ expectations of explainability. This method will enable us to anticipate potential challenges to translation as well as provide a baseline for future work to identify how explainability may evolve over time, in conjunction with the sustained use of ML tools.

Additionally, we aim to better understand which “explainability” methods can satisfy clinicians’ needs. To this end, we conducted an exploratory pilot study with 10 clinical stakeholders to: 1) understand the role of ‘explainability’ for clinical ML system; 2) identify the classes of explanation that best serve the end user, emphasizing the context and diversity of their needs; 3) characterize how these considerations can be addressed from an ML standpoint; 4) propose metrics for evaluating clinical explainability methods and identify gaps and research challenges in this area.

Technical Significance We identify when and how explainability can facilitate reliable dissemination of ML model predictions in clinical settings. We curate a set of concrete classes of explanations based on the identified clinicians’ needs. We evaluate these in the context of existing explainability literature to highlight gaps and research challenges for clinical ML. By highlighting clinicians’ needs for explanation as specific technical properties, we hope to channel explainable ML research intended for clinical end use to tackle some of the challenges.

Clinical Relevance Near future will likely see increased adoption of ML solutions in key areas of clinical practice. Building trust between clinicians and ML models, for higher uptake and satisfaction requires concerted and purposeful efforts from the ML community. This work attempts to pre-emptively identify clinicians’ needs, via research surveys, to recognize concrete technical challenges that may be addressed using explainable ML with the objective of building reliable ML systems, fostering trust for sustained use aimed at improving clinicians’ workflow and practice.

2. Study Design

We employed an upstream stakeholder engagement method (Corner *et al.*, 2012) to address the aforementioned questions. Upstream engagement here means that the survey is conducted prior to model implementation. Qualitative interviewing was determined to be an ideal method for the exploratory nature of this work, as surveys are insufficient for capturing the conceptual complexity of the phenomenon we want to understand. Furthermore, interviews provide an opportunity to clarify participants’ responses and their reasoning. Interviews reached saturation (Oreilly and Parker, 2013), defined in the literature as the point wherein no new information pertaining to the main theme of explainability arose during the interviews which is generally expected after approximately 10 interviews (Hennink *et al.*, 2017). We developed an interview guide (Jacob and Furgerson, 2012) to address the project goals. Conceptual development of ‘explainability’ proceeded throughout the interviews (Leech and Onwuegbuzie, 2009) to guide subsequent interview content Jacob and Furgerson (2012).

We interviewed 10 clinicians in two acute care settings – Intensive Care Unit (ICU) and Emergency Department (ED) with varying years of experience to develop notions of explainability and identify their needs towards building reliable ML systems for their respective clinical practice. These settings were chosen because both departments are areas where we foresee implementation of ML tools to support clinical care. Additionally, both departments have current experience working with either early warning/alert systems or opaque, non-ML decision support tools. The following describes the upstream stakeholder engagement method and study cohort in detail.

2.1. Cohort

We approached a convenience sample of stakeholder clinicians who were familiar with ML based clinical tools and are aware of ongoing developments in this field. Key stakeholders were defined as those who would be end-users of the technology and whose acceptance would ultimately determine successful clinical translation of predictive ML tools. Given the

diversity of target clinical tasks, for use in the ICU and ED settings, 6 ICU clinicians and 4 ED clinicians were selected for this exploratory survey. These individuals (5 male, 5 female) were both senior ($n = 5$, completed residency training) and junior ($n = 5$, currently in residency training) clinicians to ensure the spectrum of relevant experience was represented among the stakeholders. Discussions were held at The Hospital for Sick Children, Toronto, Canada¹ over a three week period at clinicians’ convenience, and were recorded to facilitate recall accuracy.

2.2. Procedures

Clinicians’ views were solicited to assess the concept of ML explainability with respect to their medical setting. Each interview lasted approximately 45-60 minutes. Our interviews began by exploring clinician’s notions of ‘explainability’ to determine what each clinician understood by the term and what he/she expected from ML models in their specific clinical setting. We then introduced the clinician to a hypothetical interactive scenario based on their specialties representing an ML tool incorporated into their division. For the ICU, we introduced a cardiac arrest (CA) prediction tool that reports risk of upcoming cardiac arrest (Tonekaboni *et al.*, 2018). For ED, we introduce a tool that predicts an acuity score based on triage reports (Kwon *et al.*, 2018). We queried the range of possible responses that may arise when considering different aspects of model behavior.

Hypothetical Scenarios

ICU: Suppose a machine learning-based model has been integrated into the EHR system of the hospital. As vital measurements are being constantly recorded in the ICU, this model predicts the probability of a patient having a cardiac arrest within the next few minutes, based on temporal vital measurements. The risk score is updated every few minutes based on the patients current condition. Let us say a patient in your care is being monitored by the model and you receive an alarm that notifies you that they are at imminent risk of a cardiac arrest, but you had not suspected CA for a given patient at the moment.

ED: A machine learning-based model has been integrated into the Emergency Department in your hospital. This system ranks incoming patients based on acuity using the information from the first triage. A patient in your care appears generally well, but the acuity score puts them in a high risk category.

To enhance our understanding of the concept of explainability and trust as perceived by clinicians, further questions were asked to elucidate the following constructs: how would the clinician go about evaluating the validity of an alert in terms of translating it into clinically actionable information on the part of the patient?; how does this evaluation compare to other common non-ML tools, such as alerts, risk scores, or other notifications concerning a patient’s status?; what do you currently do when one of these alerts/scores/notifications tells you something you weren’t expecting about a patient?; what information would you want about this system? Additional questions are detailed in the Appendix.

1. <http://www.sickkids.ca/>

Ethics Statement Our initial exploratory investigation aligned with the Tri-Council Policy Statement’s (TCPS2) (governing research-related activities in Canada) articles 6.11 and 10.1 stipulating activities exempt from Research Ethics Board approval. Activities exempted include: “feasibility of the research, establish research partnerships, or the design of a research proposal”.

3. Results

We organize the results and analyses of the study as follows. We identify concrete and measurable set of explanation classes curated from the qualitative assessments of the surveys (Section 3.1). We identify a few metrics that help assess the utility of any class of explanation towards its clinical utility in Section 3.3. Finally, we summarize how existing explainable ML literature fares in the context of these clinical asks. A full synopsis of the qualitative data including illustrative quotes is provided in the Appendix.

3.1. What makes a model explainable for clinicians?

Throughout the exploratory interviews it was clear that clinicians viewed explainability as a means of *justifying* their clinical decision-making (for instance, to patients and colleagues) in the context of the model’s prediction. To provide these explanations all clinicians expressed the need to understand the clinically relevant model features that align with current evidence-based medical practice. The implemented system/model needs to provide clinicians with information about the context within which the model operates and promote awareness of situations where the model may fall short (e.g., model did not use specific history or did not have information around certain aspect of a patient). *Models that fall short in accuracy were deemed acceptable so long as there is clarity around why the model under-performs.* While it is ideal to learn models based on as much contextual information, it is not always practical and clinicians are aware of such a possibility. Furthermore, the complexity of clinical medicine is such that no model is likely to achieve perfect prediction; clinicians, in fact, expect this, and the acknowledgement of this challenge promotes trustworthiness. Clearly specifying features that go into the models for decision making, is a way to facilitate trust in the model as well as directing use in specific patient populations and determining parameters guiding appropriate use. The relevant quote – ‘the variables that have derived the decision of the model’ was brought up by 3 ICU, 1 ED clinicians. This type of transparency is also identified in Mitchell *et al.* (2018) as being critical and needs to be disclosed to ML model users. While familiar metrics such as reliability, specificity, and sensitivity were important to the initial uptake of an AI tool, a critical factor for continued usage was whether the tool was *repeatedly successful in prognosticating their patient’s condition in their personal experience.* Real-world application was crucial to developing “a sense of when it’s working and when it’s limited” which meant “alignment with expectations and clinical presentation” [all clinicians].

Clinical thought process for acting on predictions of any assistive tool appears to consist of two primary steps following presentation of the model’s prediction: i) understanding, and ii) rationalizing the predictions. Thus classes of explanations for clinical ML models should be designed with the purpose of facilitating the understanding and rationalization process. Clinicians believe that carefully designed visualization and presentation can facilitate further

understanding of the model. These features are essential to sustained model use largely due to the immediacy of the clinical picture being captured and in the context of multiple competing attentional demands that require user-friendly visualization. We discerned that there are situations where well-designed visualization is not enough and only additional explanation can facilitate and fill the gap in the identified clinical workflow.

3.2. How Does this Translate to Reliable Clinical ML Design?

We determine from our discussions that a well designed explanation should augment or supplement clinical ML systems to:

- (a) Recalibrate clinician (stakeholder) trust of model predictions.
- (b) Provide a level of transparency that allows users to validate model outputs with domain knowledge.
- (c) Reliably disseminate model prediction using task specific representations (e.g. confidence scores).
- (d) Provide parsimonious and actionable steps clinicians can undertake. (e.g. potential interventions or data collection).

We further curate the following classes of explanation from qualitative assessments that clinicians identified to most effectively complement model predictions based on our survey. We highlight wherever necessary, the applicability and importance of each explanation class for different settings (ICU vs ED).

Feature Importance: Clinicians repeatedly identified that knowing the subset of features deriving the model outcome, is crucial. This allows them to compare model decision to their clinical judgment, especially in case of a discrepancy. In time-constrained settings such as the ED, important features are perceived as a crucial metric to draw the attention of clinicians to specific patient characteristics to determine how to proceed. While multiple clinicians mentioned the need to know relevant variables driving a prediction (see Appendix B), a junior clinician mentioned “you have just a number, you can still use it but in your mind when you put all the variables that make you take a decision, the weight of that variable is going to be less than if you do understand exactly what that number means”. It is also important to note that clinicians expect to see *patient specific* variable importance as well as population level variable importance (James *et al.*, 2013; Tibshirani, 1996). While an extensive survey of feature importance is beyond the scope of this work (see Table 1 for a summary), we highlight that patient-level feature importance is a challenging and far less explored machine learning problem.

Instance Level Explanations: Among commonly researched explainability methods, we investigated as part of our interviews whether data instances as explanations (Kim *et al.*, 2016; Koh and Liang, 2017) are useful in any clinical setting. Clinicians view this as finding similar patients (Sun *et al.*, 2012; Zhu *et al.*, 2016; Sharafoddini *et al.*, 2017) and believe that this kind of explanation can be only helpful in specific applications. For example, in cases where an ML model is helping clinicians find a diagnosis for a patient, it is valuable to know

the samples the model has previously seen [Cai *et al.* \(2019\)](#). However, in time constrained settings such as the ICU or ED, clinicians *did not* find this explanation technique appealing. A non-trivial challenge with this class of explanation, that researchers need to be aware of, is the definition of similarity that is used ([Santini and Jain, 1999](#)). Clinicians identified that despite similar outcomes patients may differ significantly in the clinical trajectory to arrive of those outcomes, and vice versa. Various definitions for similarity may be proposed depending on the task the model is performing, and the way the user (clinician) wants to use the explanation. For example, some clinicians were interested in seeing similar samples to their patients because they were curious to know what actions were taken in those cases, and the associated outcomes of the interventions.

Uncertainty: Clinicians overwhelmingly indicated that the model’s overall accuracy was not sufficient and its clinical alignment in their judgment often determined their sustained use and trust in the model. That is to say that each time an alert presents a prediction, clinicians will anticipate a clinically significant change in the patient’s status that should ideally align with the model’s prediction. Previous reports have indicated that despite the determination of expert-agreed threshold for triggering alerts, there are still many alerts not accompanied by a clinically actionable change ([Umscheid *et al.*, 2015](#)) which undermines use and endorsement by clinicians ([Guidi *et al.*, 2015](#)). This point holds particular importance for clinical ML to ensure developments align with clinicians’ expectations to promote trust and sustained model use. Presenting certainty score on model performance or predictions is perceived by clinicians as a sort of explanation that complements the output result. They also suggested that this score can be used as a threshold for reporting results only when they imply that model is very certain of its prediction. Many clinicians noted that “alarm or click fatigue” (indicating the annoyance with repeating response prompts through the EHR system) ([Embi and Leonard, 2012](#)) is a significant concern that may be worsened by prediction tools. This issue is ubiquitous across healthcare contexts and requires careful consideration by model developers so as to not perpetuate clinician disillusionment and disengagement with these systems. This makes calibration of complex models ([Guo *et al.*, 2017](#)) a significant technical challenge that needs to be addressed for clinical practice.

An additional challenge is the fact that even models performing acceptably on average can have significant individual level errors ([Nushi *et al.*, 2018](#)) which is undesirable in clinical practice. There are two sources for uncertainty we highlight ([Gal, 2016](#)) that can affect model trust:

1. Model Uncertainty: Uncertainty of model parameters that best explain observed data is known to be a significant source of uncertainty in model performance ([Nuzzo, 2014](#); [Zhang *et al.*, 2013](#)). Modeling for this uncertainty can be achieved by actively accounting for distributional differences during training ([Gal and Ghahramani, 2016](#); [Schulam and Saria, 2019](#)). Few of these methods have been actively co-opted in clinical healthcare for explainability but nonetheless are crucial to uptake.

In addition, model mis-specification refers to performance deterioration due to a mis-specification of the model class used during training. It can lead to an overall under-performing system and can further reduce the effectiveness of some methods used for dealing with model uncertainty ([Wen *et al.*, 2014](#)). This problem is especially understudied for time-series models that heavily employ deep learning methods ([Zhu and](#)

Laptev, 2017) and to the best of our knowledge remains mostly unexplored in clinical healthcare employing deep learning methods.

2. **Data Uncertainty:** This type of uncertainty can come from noisy, missing data or an existing inherent uncertainty in the data. Both of these challenges are commonly faced in clinical ML. Inherently uncertain observation can be for instance a complicated patient with difficult diagnosis, thus is considered a relatively more difficult problem. Characterizing consistency under missingness has been only briefly studied in supervised classification system (Josse *et al.*, 2019) and needs to be rigorously adopted and evaluated for clinical applications.

Temporal Explanations: ‘Patient trajectories that are influential in driving model predictions’ was reported to be an important aspect of model explanation. This area remains relatively unexplored in clinical ML explainability literature with a few exceptions (Xu *et al.*, 2018; Choi *et al.*, 2016; Yang *et al.*, 2018) focused primarily on attention based and sensitivity analysis mechanisms in deep learning (Vaswani *et al.*, 2017). However, explanations based on attention mechanisms can produce inconsistent explanations (Jain and Wallace, 2019). Explainability for temporal data is a relatively challenging task as it requires analyzing higher order temporal dependencies. ML models that perform such predictions, should be able to explain their prediction based on changes in individual patient state. In units such as the ICU, clinicians are interested to see the change of state that has resulted in a certain prediction. Note that this is also related to explanations using Feature Importance (see Sub-section 3.2) as a necessary component of providing explanations in the temporal domain. To the best of our knowledge, not much has been investigated in this domain, as formal definitions of explanations.

Transparent Design Clinicians pointed to the need for models that reflect a similar analytic process to the established methodology of evidence-based medical decision making (Haynes *et al.*, 1997). They anticipate model decisions with similar emphasis on recognizing the clinical features driving the prediction as being more interpretable since it reflects how they currently assess a patient’s risk status. Translating a model to a transparent design (e.g. a decision tree) is useful to facilitate rationalization of model behavior, as expressed by one ICU clinician: “would want to know the equation to know what the weights are - but if it is variable and if knowing that is too much detail then its just not that helpful.” [1 Senior ICU clinician] and “if there is a discrepancy between the current clinical protocol and the model, and we dont know why this is occurring, we are going to be nervous” [1 Senior ED clinician]. To the best of our knowledge, this is the most well studied area of model explanations in general and clinical ML literature. For example, a variety of rule based methods have been proposed precisely for the purpose of transparent clinical design (Lakkaraju *et al.*, 2016; Wang and Rudin, 2015). Model distillation (Hinton *et al.*, 2015) refers to learning a “simpler” model class that perform at par with an complex ML model and can be considered as an attempt toward transparent design (Che *et al.*, 2015). Closely related are regularization based methods (Wu *et al.*, 2018) that attempt to regularize models for enhanced interpretability. These methods tend to assume there are trade-offs between model performance and their perceived “explainability”. It is unclear

when such a trade-off is acceptable, especially if it amplifies individual level errors. We further argue that such models should be rigorously evaluated for performance under other factors like mis-specification of the distilled model or regularized model, induced biases and general uncertainty to facilitate reliable adoption to clinical practice.

3.3. Metrics for Evaluating Explanations

We further identify that explanations intended to build trust in clinical settings could benefit from being rigorously evaluated against the following metrics:

Domain Appropriate Representation The quality of the explanation should be evaluated in terms of whether the representation is coherent with respect to the application task. For instance, for patients in the ICU, clinicians already have a lot of context as to why the patient is currently admitted to this division. Explanations that are redundant to the ICU task are not desirable unless critical to potential clinical workflows. The interviews revealed a diversity of settings in which clinicians operate, like diagnoses versus cardiac care management (see Instance based explanations in Section 3.1) which determines the representation best suited for appropriate settings. The representation itself should not further obfuscate model behavior for the clinician. This for instance requires careful filtering of the information most useful at any instant (e.g. highlighting “age” even if it is an important feature is not useful if all potential clinical followups are uniform across age or it may not be helpful to receive an alert for a patient who is already en route to receiving a known intervention, like a surgery). The type of explanation will largely be evaluated by rigorous user studies involving stakeholders (Doshi-Velez and Kim, 2017).

Potential Actionability Given a model that satisfies minimal trust based criterion and has been sufficiently evaluated for performance, any complementary explanation should inform follow-up clinical workflow, including rationalization of model prediction (see Section 3.1). This followup can be anything ranging from checking on the patient to ordering additional lab measurements or determining an intervention. Overall, in an applied field like clinical ML, explanations that are informative (in terms of understanding the mode), but have no impact on the workflow are of less importance. Similarly, the explanation should be parsimonious and timely. The nature of explanation should also account for inter-dependant factors that inform its usability. This can be determined by the purpose it will serve in any specific setting. For instance, ED clinicians believed that any risk assesment tool will help them allocate resources appropriately, therefore the explanations provided for the ML model should facilitate the decision making. For instance, in an ICU setting, when time to action is extremely limited, counterfactual explanations or patient similarity are not informative, usable, and result in cognitive overload (Henriksen *et al.*, 2008).

Consistency Consistency refers to two factors: (i) The set of explanations should be injective (or changes in model predictions should yield discernable changes in the explanation) (ii) These changes should be invariant to underlying design variations and should only reflect relevant clinical variability. Explanations that are inconsistent across any of these factors effectively violate their reliable actionability, and also negatively impact the trust of clinicians. Consistency is also closely related to robustness of explanations. For instance, such lack of consistency has already been identified using statistical methods for

commonly used explanation techniques in deep learning literature (Adebayo *et al.*, 2018; Jain and Wallace, 2019) outside of clinical ML.

3.4. What is Missing in Explainable ML?

While a multitude of novel explainability methods have been proposed in the literature (Guidotti *et al.*, 2018), they may not be directly applicable or have to be significantly tailored to clinical settings in the context of asks outlined in Section 3.1. In the following, we focus on identifying a collection of existing state of the art methods than can be used to generate explanations that are usable for clinical stakeholders. We summarize our results in Table 1 and identify possible shortcomings of existing ML approaches for clinical usability and relevance. Note that while a complete review of existing explainability methods is outside the scope of this work, we highlight the most relevant and popular techniques in the context of the classes outlined above.

Note that even if an existing explainability method is evaluated in a clinical setting, a rigorous evaluation of these methods for robustness and other inter-dependant factors still necessitates sustained research. Additionally, as noted in Section 4 and Table 1, some methods may work well for specific kinds of data, like images and may have to be non-trivially extended to be applicable to other data, and within the additional complexity of clinical settings.

4. Related Work

To the best of our knowledge, no prior works have conducted target stakeholder studies for the specific purpose of identifying explainability challenges in clinical ML. In terms of translating ML methods to clinical practice, Escobar *et al.* (2016); Elish (2018) have extensively recorded their observations in piloting an early sepsis warning system to a single institute over a period of multiple days. A few user studies have evaluated prototypical Electronic Health Recording systems (Mazur *et al.*, 2016) or associated visualization prototypes (Ghassemi *et al.*, 2018a) for efficacy of facilitating diverse clinical workflows (Nolan *et al.*, 2017). Herein we highlight some state of the art literature in explainable ML which we have discussed in terms of their relevance to adoption of ML models in clinical practice in Section 3.1.

Explainability in general machine learning has focused on understanding model behavior from different perspectives. Although simpler model classes (Lundberg *et al.*, 2018; Wu *et al.*, 2018) are considered more interpretable by users, a general trade-off between model performance and explainability is assumed (Caruana *et al.*, 2015). Some methods like LIME and Anchors (Ribeiro *et al.*, 2016, 2018) use simpler model classes to derive feature level explanations (Sundararajan *et al.*, 2017). Abstractions such as data instances (Koh and Liang, 2017; Kim *et al.*, 2016) are used as a means of explaining model behaviors. Visualization is also explored as a method of providing an insight in model behaviour (Yosinski *et al.*, 2015; Ming *et al.*, 2017). We characterize these methods as *diagnostic* methods as they are not necessarily targeted to provide explanations for the target user but focus on improving model understanding. See Guidotti *et al.* (2018) for an extensive survey of existing methods for explainable ML in general.

Table 1: Summary of Explainable ML methods Contextualized for Clinical Applicability

Explanation Class	Representative Existing Methods	Possible shortcomings for clinical settings
Feature Importance	Sensitivity Analysis (Saltelli et al., 2008), LRP (Bach et al., 2015) (Yang et al., 2018) LIME, Anchors, Shapley Values (Ribeiro et al., 2016, 2018 ; Lundberg et al., 2018)	Complex correlation between features of clinical models can be a challenge Further evaluation for consistency required
Instance Level Explanation	Influence functions (Koh and Liang, 2017) Prototypes and Criticisms (Kim et al., 2016)	Not evaluated on complex clinical models Limited applicability
Uncertainty	Distributional shift (Subbaswamy and Saria, 2018) Parameter uncertainty (Gal and Ghahramani, 2016 ; Schulam and Saria, 2019)	Not evaluated on complex clinical models
Temporal Explanations	RETAIN, RAIM (Xu et al., 2018 ; Choi et al., 2016)	Potential lack of consistency due to the attention mechanism (Jain and Wallace, 2019)
Transparent Design	Rule Based Methods (Lakkaraju et al., 2016 ; Wang and Rudin, 2015)	Less powerful in modeling more complex applications; Generally assume a trade-off of accuracy and explainability

In clinical ML, visualization systems have been evaluated for efficacy of interaction between EHR systems and healthcare practitioners ([Ghassemi et al., 2018a](#)). Neural attention mechanisms ([Xu et al., 2015](#)) are used for developing interpretable models for heart failure and cataract on-set prediction ([Kwon et al., 2019](#); [Choi et al., 2016](#)). [Wu et al. \(2018\)](#) use tree-based regularization for sepsis and HIV treatment, while [Wang and Rudin \(2015\)](#); [Ustun and Rudin \(2016\)](#) build rule based or sparse linear methods for predicting hospital readmissions and sleep apnea scoring systems. [Ahmad et al. \(2018\)](#) conduct the most extensive survey of explainable methods in clinical healthcare, and analyze them in the context of existing notions in general ML and how they relate to clinical healthcare, discussing ethical questions thereof. They conclude that general agreement in the field of clinical healthcare toward explainability is fairly low. Our work attempts to bridge some of this uncertainty by initiating a conversation with and accounting for the needs of stakeholders.

Several works have attempted to codify rigorous evaluation of explainability methods. For instance, [Doshi-Velez and Kim \(2017\)](#) propose ways to evaluate interpretability methods

depending on application type and thereby the nature of evaluation, while [Narayanan *et al.* \(2018\)](#) and [Poursabzi-Sangdeh *et al.* \(2018\)](#) conduct user studies to understand the quality of many proposed explainability methods. Statistical methods have also been introduced for evaluating the robustness of a few candidate explanation methods like saliency and neural attention architectures ([Adebayo *et al.*, 2018](#); [Jain and Wallace, 2019](#)).

5. Discussion

This work documents the ongoing challenge of translation of clinical ML with a particular focus on explainability through the eyes of end users. We demonstrate how clinicians’ views sometimes differ from existing notions of explainability in ML, and propose strategies for enhancing buy-in and trust by focusing on these needs. In light of the objectives highlighted in Section 1, we survey clinicians with diverse specialties and identify when explainability methods can assist in enhancing clinicians’ trust in ML models. Our research survey involved creating hypothetical scenarios of deploying a machine learning based predictive tool to carry out specific tasks in the ICU and the ED respectively. We demonstrate that by accounting for target stakeholders, even though the explainability task in clinical settings is significantly diverse, but it can be codified into specific technical challenges. We further outline the need to evaluate clinical explainability methods rigorously under the proposed metrics in Section 3.3. To the best of our knowledge, this is the first attempt at involving ICU and ED stakeholders to identify targeted clinical needs and evaluating it against general machine learning literature in this field.

Some of our non-technical observations, as highlighted in the qualitative synopsis in Section 3.1 and Appendix B are corroborated to some extent by the observations of [Elish \(2018\)](#) who followed the development and deployment of a machine learning based sepsis risk detection tool in a clinical setting. Our work however is far more general, as we concretely map conclusions from clinical surveys to prominent gaps in explainable ML literature as it pertains to effective clinical practice. Limitations of our methods are highlighted below.

5.1. Limitations and Future work

Our research survey for identifying explainability challenges was restricted to ICU and ED specialists. Most specialists we interviewed have reasonable knowledge of clinical ML systems (aware of or involved in academic research involving clinical ML systems). Given the research questions which were focused on prediction tasks, the identified challenges may only have limited applicability to wider class of ML models (for example, reinforcement learning, survival models etc). Applicability of the identified classes of explanations to other clinical settings (like outpatient tasks) needs further evaluation. While our study was an exploration into the research design of a broader study of explainability, we were able to glean useful insights.

A strength of the study is that we attempted to elucidate concepts under the explainability framework that can apply to many novel ML-based implementations. A larger research question we are currently pursuing is whether satisfaction of these explainability requirements is directly related to perceived usefulness of the tool as well as to sustained use and endorsement. We hope to develop a conceptual framework that can be utilized across

a spectrum of ML tools that can support consistency and quality of ML’s explainability requirements in hospital environments.

Acknowledgments

We would like to thank all the clinicians who volunteered their valuable time for this exploratory study.

References

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. (2018). Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pages 9525–9536.
- Ahmad, M. A., Eckert, C., and Teredesai, A. (2018). Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 559–560. ACM.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, **10**.
- Bedoya, A. D., Clement, M. E., Phelan, M., Steorts, R. C., O'Brien, C., and Goldstein, B. A. (2019). Minimal impact of implemented early warning score and best practice alert for patient deterioration. *Critical care medicine*, **47**.
- Cai, C. J., Reif, E., Hegde, N., Hipp, J., Kim, B., Smilkov, D., Wattenberg, M., Viegas, F., Corrado, G. S., Stumpe, M. C., *et al.* (2019). Human-centered tools for coping with imperfect algorithms during medical decision-making. *arXiv preprint arXiv:1902.02960*.
- Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., and Elhadad, N. (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1721–1730. ACM.
- Che, Z., Purushotham, S., Khemani, R., and Liu, Y. (2015). Distilling knowledge from deep networks with applications to healthcare domain. *arXiv preprint arXiv:1512.03542*.
- Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., and Stewart, W. (2016). Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems*, pages 3504–3512.
- Corner, A., Pidgeon, N., and Parkhill, K. (2012). Perceptions of geoengineering: public attitudes, stakeholder perspectives, and the challenge of upstream engagement. *Wiley Interdisciplinary Reviews: Climate Change*, **3**(5), 451–466.
- Doshi-Velez, F. and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

- Duncan, H., Hutchison, J., and Parshuram, C. S. (2006). The pediatric early warning system score: a severity of illness score to predict urgent medical need in hospitalized children. *Journal of critical care*, **21**(3), 271–278.
- Elish, M. C. (2018). The stakes of uncertainty: Developing and integrating machine learning in clinical care. In *Ethnographic Praxis in Industry Conference Proceedings*, volume 2018, pages 364–380. Wiley Online Library.
- Embi, P. J. and Leonard, A. C. (2012). Evaluating alert fatigue over time to ehr-based clinical trial alerts: findings from a randomized controlled study. *Journal of the American Medical Informatics Association*, **19**(e1), e145–e148.
- Escobar, G. J., Turk, B. J., Ragins, A., Ha, J., Hoberman, B., LeVine, S. M., Balleca, M. A., Liu, V., and Kipnis, P. (2016). Piloting electronic medical record–based early detection of inpatient deterioration in community hospitals. *Journal of hospital medicine*, **11**.
- Gal, Y. (2016). *Uncertainty in deep learning*. Ph.D. thesis, PhD thesis, University of Cambridge.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059.
- Ghassemi, M., Pushkarna, M., Wexler, J., Johnson, J., and Varghese, P. (2018a). Clinicalvis: Supporting clinical task-focused design evaluation. *arXiv preprint arXiv:1810.05798*.
- Ghassemi, M., Naumann, T., Schulam, P., Beam, A. L., and Ranganath, R. (2018b). Opportunities in machine learning for healthcare. *arXiv preprint arXiv:1806.00388*.
- Guidi, J. L., Clark, K., Upton, M. T., Faust, H., Umscheid, C. A., Lane-Fall, M. B., Mikkelsen, M. E., Schweickert, W. D., Vanzandbergen, C. A., Betesh, J., *et al.* (2015). Clinician perception of the effectiveness of an automated early warning and response system for sepsis in an academic medical center. *Annals of the American Thoracic Society*, **12**.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, **51**.
- Gunning, D. (2017). Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA)*, *nd Web*.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1321–1330. JMLR. org.
- Haynes, R. B., Sackett, D. L., Richardson, W. S., Rosenberg, W., and Langley, G. R. (1997). Evidence-based medicine: How to practice & teach ebm. *Canadian Medical Association Journal*, **157**(6), 788.

- Hennink, M. M., Kaiser, B. N., and Marconi, V. C. (2017). Code saturation versus meaning saturation: how many interviews are enough? *Qualitative health research*, **27**(4), 591–608.
- Henriksen, K., Dayton, E., Keyes, M. A., Carayon, P., and Hughes, R. (2008). Understanding adverse events: a human factors framework.
- Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Jacob, S. A. and Furgerson, S. P. (2012). Writing interview protocols and conducting interviews: Tips for students new to the field of qualitative research. *The qualitative report*, **17**(42), 1–10.
- Jain, S. and Wallace, B. C. (2019). Attention is not explanation.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An introduction to statistical learning*, volume 112. Springer.
- Josse, J., Prost, N., Scornet, E., and Varoquaux, G. (2019). On the consistency of supervised learning with missing values. *arXiv preprint arXiv:1902.06931*.
- Kim, B., Khanna, R., and Koyejo, O. O. (2016). Examples are not enough, learn to criticize! criticism for interpretability. In *Advances in Neural Information Processing Systems*, pages 2280–2288.
- Koh, P. W. and Liang, P. (2017). Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1885–1894. JMLR. org.
- Kwon, B. C., Choi, M.-J., Kim, J. T., Choi, E., Kim, Y. B., Kwon, S., Sun, J., and Choo, J. (2019). Retainvis: Visual analytics with interpretable and interactive recurrent neural networks on electronic medical records. *IEEE transactions on visualization and computer graphics*, **25**.
- Kwon, J.-m., Lee, Y., Lee, Y., Lee, S., Park, H., and Park, J. (2018). Validation of deep-learning-based triage and acuity score using a large national dataset. *PloS one*, **13**(10), e0205836.
- Lakkaraju, H., Bach, S. H., and Leskovec, J. (2016). Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, New York, NY, USA. ACM.
- Leech, N. L. and Onwuegbuzie, A. J. (2009). A typology of mixed methods research designs. *Quality & quantity*, **43**.
- Lipton, Z. C. (2016). The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*.

- Lundberg, S. M., Nair, B., Vavilala, M. S., Horibe, M., Eisses, M. J., Adams, T., Liston, D. E., Low, D. K.-W., Newman, S.-F., Kim, J., *et al.* (2018). Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature biomedical engineering*, **2**(10), 749.
- Masino, A. J., Harris, M. C., Forsyth, D., Ostapenko, S., Srinivasan, L., Bonafide, C. P., Balamuth, F., Schmatz, M., and Grundmeier, R. W. (2019). Machine learning models for early sepsis recognition in the neonatal intensive care unit using readily available electronic health record data. *PloS one*, **14**.
- Mazur, L. M., Mosaly, P. R., Moore, C., Comitz, E., Yu, F., Falchook, A. D., Eblan, M. J., Hoyle, L. M., Tracton, G., Chera, B. S., *et al.* (2016). Toward a better understanding of task demands, workload, and performance during physician-computer interactions. *Journal of the American Medical Informatics Association*, **23**(6), 1113–1120.
- Miller, T. (2018). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*.
- Ming, Y., Cao, S., Zhang, R., Li, Z., Chen, Y., Song, Y., and Qu, H. (2017). Understanding hidden memories of recurrent neural networks. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 13–24. IEEE.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., and Gebru, T. (2018). Model cards for model reporting. *arXiv preprint arXiv:1810.03993*.
- Narayanan, M., Chen, E., He, J., Kim, B., Gershman, S., and Doshi-Velez, F. (2018). How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682*.
- Nolan, M. E., Cartin-Ceba, R., Moreno-Franco, P., Pickering, B., and Herasevich, V. (2017). A multisite survey study of emr review habits, information needs, and display preferences among medical icu clinicians evaluating new patients. *Applied clinical informatics*, **8**(04), 1197–1207.
- Nushi, B., Kamar, E., and Horvitz, E. (2018). Towards accountable ai: Hybrid human-machine analyses for characterizing system failure. In *Sixth AAAI Conference on Human Computation and Crowdsourcing*.
- Nuzzo, R. (2014). Scientific method: statistical errors. *Nature News*, **506**.
- Oreilly, M. and Parker, N. (2013). unsatisfactory saturation: a critical exploration of the notion of saturated sample sizes in qualitative research. *Qualitative research*, **13**(2), 190–197.
- Páez, A. (2019). The pragmatic turn in explainable artificial intelligence (xai). *Minds and Machines*, pages 1–19.

- Papernot, N., McDaniel, P., and Goodfellow, I. (2016). Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*.
- Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Vaughan, J. W., and Wallach, H. (2018). Manipulating and measuring model interpretability. *arXiv preprint arXiv:1802.07810*.
- Prytherch, D. R., Smith, G. B., Schmidt, P. E., and Featherstone, P. I. (2010). Viewstowards a national early warning score for detecting adult inpatient deterioration. *Resuscitation*, **81**(8), 932–937.
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, New York, NY, USA. ACM.
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S. (2008). *Global sensitivity analysis: the primer*. John Wiley & Sons.
- Santini, S. and Jain, R. (1999). Similarity measures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**.
- Schulam, P. and Saria, S. (2019). Auditing pointwise reliability subsequent to training. *arXiv preprint arXiv:1901.00403*.
- Sharafoddini, A., Dubin, J. A., and Lee, J. (2017). Patient similarity in prediction models based on health data: a scoping review. *JMIR medical informatics*, **5**.
- Smith, G. B., Prytherch, D. R., Meredith, P., Schmidt, P. E., and Featherstone, P. I. (2013). The ability of the national early warning score (news) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death. *Resuscitation*, **84**.
- Subbaswamy, A. and Saria, S. (2018). Counterfactual normalization: proactively addressing dataset shift using causal mechanisms. In *Uncertainty in Artificial Intelligence*, pages 947–957.
- Sun, J., Wang, F., Hu, J., and Edabollahi, S. (2012). Supervised patient similarity measure of heterogeneous patient records. *ACM SIGKDD Explorations Newsletter*, **14**.
- Sundararajan, M., Taly, A., and Yan, Q. (2017). Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3319–3328. JMLR. org.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, **58**.

- Tonekaboni, S., Mazwi, M., Laussen, P., Eytan, D., Greer, R., Goodfellow, S. D., Goodwin, A., Brudno, M., and Goldenberg, A. (2018). Prediction of cardiac arrest from physiological signals in the pediatric icu. In *Machine Learning for Healthcare Conference*, pages 534–550.
- Umscheid, C. A., Betesh, J., VanZandbergen, C., Hanish, A., Tait, G., Mikkelsen, M. E., French, B., and Fuchs, B. D. (2015). Development, implementation, and impact of an automated early warning and response system for sepsis. *Journal of hospital medicine*.
- Ustun, B. and Rudin, C. (2016). Supersparse linear integer models for optimized medical scoring systems. *Machine Learning*, **102**.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Vayena, E., Blasimme, A., and Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS medicine*, **15**.
- Vincent, J.-L., Moreno, R., Takala, J., Willatts, S., De Mendonça, A., Bruining, H., Reinhart, C., Suter, P., and Thijs, L. (1996). The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. *Intensive care medicine*, **22**(7), 707–710.
- Wang, F. and Rudin, C. (2015). Falling rule lists. In *Artificial Intelligence and Statistics*, pages 1013–1022.
- Wen, J., Yu, C.-N., and Greiner, R. (2014). Robust learning under uncertain test distributions: Relating covariate shift to model misspecification. In *ICML*, pages 631–639.
- Wu, M., Hughes, M. C., Parbhoo, S., Zazzi, M., Roth, V., and Doshi-Velez, F. (2018). Beyond sparsity: Tree regularization of deep models for interpretability. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057.
- Xu, Y., Biswal, S., Deshpande, S. R., Maher, K. O., and Sun, J. (2018). Raim: Recurrent attentive and intensive model of multimodal patient monitoring data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2565–2573. ACM.
- Yang, Y., Tresp, V., Wunderle, M., and Fasching, P. A. (2018). Explaining therapy predictions with layer-wise relevance propagation in neural networks. In *2018 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 152–162. IEEE.
- Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., and Lipson, H. (2015). Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*.

- Zhang, K., Schölkopf, B., Muandet, K., and Wang, Z. (2013). Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pages 819–827.
- Zhu, L. and Laptev, N. (2017). Deep and confident prediction for time series at uber. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 103–110. IEEE.
- Zhu, Z., Yin, C., Qian, B., Cheng, Y., Wei, J., and Wang, F. (2016). Measuring patient similarities via a deep architecture with medical concept embedding. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*.

Appendix A. Interview Protocol

Additional Questions

These questions were asked to get clinicians’ perspective regarding specific explainability features (common in existing explainable ML literature as curated by the authors). We additionally asked questions that allowed us to get a sense of potential metrics clinicians trust and would like the model to be supplemented with.

1. What kind of information could help you feel confident in how you chose to manage a patient?
2. What kinds of things do you look to now, for example, to help you estimate how a patient will do?
3. If you know the confidence of the models prediction for a specific instance, would that make any difference in your reaction?
4. If the model presented to you the part of the patient measurements and history that has influenced the model decision the most, would you find it helpful? Do you think it could help save time when making a decision on interventions?
5. Do you think it would be helpful to get examples of similar patient with the similar trajectory? If so, what would a similar case look like for you? How would you use this information? Do you want to see similarity of clinical outcome or trajectory?
6. Are the set of top measurements that are driving the prediction useful?
7. Given all the points that were brought up (enumerate all options) , if you were back in the ICU/ED scenario, which one of the above added information would you find more useful and actionable?

Appendix B. Qualitative Synopsis and Representative Quotes from Stakeholders

Here we include representative quotes from stakeholders as evidence that formed the basis of our results. We present these in terms of general perception of clinicians towards ML systems as well as their perception of explainability.

General Perception of the Need for Clinical ML Systems: Clinical ML systems that could support the improvement of patient outcomes were viewed as an important adjunct to facilitate care in both ICU and ED settings. Specific needs identified included the ability to have a system as a constant surveillance (e.g., collecting and analyzing physiologic signals), which was likened to having “an extra tool in [their] toolbox” or a “failsafe” [3 ED, 2 ICU, mostly junior], i.e. a tool that will be constantly monitoring the situation and alerting clinicians to the onset of potentially critical events. Several clinicians noted that the ability to intervene earlier was important in order to save time or prevent or minimize bad outcomes (e.g., loss of functionality, death). Successful ML implementations in this regard were likened to a senior clinician who “picks up every sign so well” [3 ICU, mostly junior],

in that the ML systems’ value lies in its ability to mimic the clinical acuity of experienced specialists.

Perhaps the most important benefit of ML was the ability to facilitate a risk- or acuity-based allocation of attention in the ED. Clinicians in the ED setting particularly described the challenge of not being able to predict with certainty the trajectory of individual patients, and interventions were thus necessarily reactive and subsequent to a negative clinical condition (e.g., blood pressure dropping suddenly). ML should be like having an “constant eye on the patient” [all clinicians] that could support the direction of appropriately timed resources to the bedside of a patient prior to the point of clinical decline. Supporting resource allocation at the precise moment of need was particularly beneficial in the ED, where “there are a thousand things going on at once and you have incomplete information” [all ED clinicians]. In supporting appropriate attention allocation, clinicians refer to the ability of the ML model to provide a quick, reliable piece of information that clinicians can trust to direct them to the right patient at the right time.

Clinical Perception of Explainability of ML Systems: Primarily, clinicians viewed explainability as a means of *justifying* their clinical decision-making (for instance, to patients and colleagues) in the context of the model’s prediction. To them, the term often implied awareness of “the variables that have derived the decision of the model” [3 ICU, 1 ED]. One senior clinician noted, “if it’s on a computer then the assumption is that somebody had already figured it out,” implying that a lot of features may not require lengthy explanations but instead would require clinical accuracy, as we described in Section 3.1, (Uncertainty). Explanations seemed to be closely tied to visualization and representation of model predictions. A few mentioned it might be helpful to know “parameters that are feeding the model... what data is processed and aggregated” [1 Junior ICU clinician]. Clinicians also noted “if there is a discrepancy between the current clinical protocol and the model, and we don’t know why this is occurring, we are going to be nervous” [1 Senior ED clinician]. Universally, awareness of the factors driving the prediction were viewed as essential to determining whether to do additional tests, return to bedside, or intervene: “you have just a number, you can still use it but in your mind when you put all the variables that make you take a decision, the weight of that variable is going to be less than if you do understand exactly what that number means” [1 Junior ICU clinician]. Lastly, for sustained improvement using clinical feedback, knowledge of the model’s development was perceived as important insofar as it facilitated the refinement of the tool itself. For example, one clinician noted it is “crucial for clinicians to understand how the model is making predictions to be able to provide feedback on specific conditions that can influence the false positive rate” [1 Junior ED clinician].