Information Extraction from Patient Care Reports for Intelligent Emergency Medical Services

Sion Kim, Weishi Guo, Ronald Williams, John Stankovic, Homa Alemzadeh
University of Virginia, Charlottesville, VA 22904, USA
{sk9uth, wg7qn, rdw, jas9f, ha4d}@virginia.edu

***Abstract*—Every year, more than 30 million emergency medical incidents are responded to in the U.S. Upon arrival at an incident scene, responders assess the situation and provide emergency medical care to the patients before transporting them to hospitals. In this process, the responders collect substantial amounts of data with different levels of importance and confidence, including the patient's present medical conditions, past medical history, and interventions performed. Although there are several standards and tools for collecting, storing, and sharing EMS data, less attention has been given to reliably translating this wealth of information into actionable knowledge for assessing the performance of emergency operations and evaluating response protocols. This paper presents the analysis of over 35,900 EMS pre-hospital electronic Patient Care Reports (ePCR) from an urban ambulance agency. We used both the structured and unstructured information in the dataset to develop a domain-specific EMS ontology with a standardized lexicon for medications, procedures, responders' impressions, call types, chief complaints, and signs and symptoms. The EMS ontology was used to develop methods for automated segmentation of narratives, detection and correction of incorrect/incomplete information in the reports, and generation of time-series data to represent the progression of incidents and the most common sequences of response actions (models of EMS protocols). Finally, we performed an analysis on the relationships among different aspects of incidents to provide insights for the design of future EMS assistive technologies.**

***Index Terms*—emergency medical services, electronic patient care report, EMS, ePCR.**

I. Introduction

Emergency Medical Services (EMS) responders provide both basic and advanced life support for serious and life-threatening medical conditions and traumatic injuries, with the means of stabilizing the patients and providing transport to definitive care. These responders are trained to follow established emergency response protocols based on their assessment of the situation. To assess the situation, responders work with a substantial amount of information, typically considering circumstances and history of the incident, illness, or injury including patient symptoms and medical history. Additional data and direction are collected through real-time sensor measurements from medical devices, voice communications with medical command, and discussions with bystanders or family members. Filtering, processing, and recording information with different levels of importance and confidence require a significant amount of responders' cognitive effort.

This work was partially supported by the award 60NANB17D162 from the U.S. Department of Commerce, National Institute of Standards and Technology (NIST), and a research grant from the Commonwealth Cyber Initiative (CCI).

In addition to the cognitive demands, there are challenges in making sense of the substantial amount of EMS data that is documented manually by responders. State and national frameworks for data collection such as the National EMS Information System (NEMSIS) exist [1]. However, even with some level of standardization, the collected data is often unreliable, incomplete, inconsistent, and might be incorrectly entered into the reports [2]. This is due to the time pressure and stresses the responders must deal with when documenting their observations and actions concurrently with the unfolding incident. There exist tools and devices to digitally document incidents, such as ImageTrend [3], Emergency Department Information Exchange [4], and GRACE [5]. However, most of these technologies require initial manual input, which is error-prone or not robust enough to the noise in the environment. Further, a significant part of EMS data are narratives described in free-form natural language text, which are difficult to analyze without understanding the domain-specific semantics, terminology, and abbreviations, and considering the incident context. If these problems are addressed, the vast amount of documented information is an untapped resource for assessing emergency operations performance and improving response protocols.

Previous works have focused on analyzing EMS data for understanding national trends, responses to specific illnesses or incident types, developing data standardization systems and assistive technologies for decision support and incident reporting, and designing Natural Language Processing (NLP) techniques for analysis of clinical text and Electronic Health Records (EHR). But there are still several challenges and research gaps that need to be addressed:

- Lack of unified labeled datasets. Most of the existing data from previously collected EMS datasets, are inconsistent, incomplete, error-prone, and often not labeled according to the EMS guidelines. A significant amount of manual effort and domain expertise is needed for labeling and making sense of such data.
- Little previous work on general analyses of the relationships among different aspects of EMS data to derive insights that can inform the development of decision support tools and assistive technologies for responders. Most previous works have focused on specific illnesses or incident types such as stroke or opioid overdose.
- Lack of standard lexicons and ontologies for EMS that can facilitate automated analysis and summarization of

incident data. The existing standard systems and frameworks for EMS mainly focus on unified data collection rather than semantic understanding and data analysis.
- Domain mismatch between EHR and ePCR. Although there are many existing analytic solutions for clinical text and EHR, they are not directly applicable to the EMS domain. This is because of the different nature and structure of pre-hospital ePCR, which are being collected over a short time frame and contain unconfirmed diagnosis and domain-specific terminology and abbreviations.

This paper presents a comprehensive and generalized study of EMS incidents based on a large dataset of pre-hospital electronic Patient Care Reports (ePCR) from an urban ambulance agency. We developed a domain-specific EMS ontology along with analytic methods for automated completion and correction of incident reports and generation of EMS protocol models that can facilitate both offline and online analysis of big EMS data and the development of future assistive technologies for responders.

The main contributions of this work include:

- Semi-automated development of a domain-specific ontology for EMS with standardized lexicon for medications, procedures, impressions, chief complaints, call types, and signs and symptoms.
- Automated segmentation of EMS narratives to extract and complete information on chief complaints, history of illness, medical history, medications, and allergies.
- Automated extraction and correction of patients' demographic information (e.g., gender, age) from narratives.
- Automated generation of incident time-series and treatment sequences for modeling EMS protocols which enables detection of discrepancies and missing information.
- Comprehensive analysis of trends and relationships among different aspects of response in EMS incidents.

The automated narrative segmentation, extraction of demographic information, and correction/completion of data can be also conducted online for more accurate EMS form filling [5] and decision support [6]. The insights from this study can provide a data-driven basis for evaluating the performance of emergency response operations and augment existing approaches to modeling EMS protocols and interventions. The code, resulting conclusions, and the EMS ontology generated from this analysis are made publicly available to the research community through an online repository[1].

II. Related Work

Cognitive Assistance: Previous works to address the cognitive overload of responders have proposed the use of assistive technologies with advanced sensing, computing, and artificial intelligence capabilities for data collection, decision support, and incident reporting. One work, CognitiveEMS [6], [7], is a cognitive assistant system that improves responders' situational awareness by automated collection and analysis of data in real-time during an incident and providing protocol-driven feedback to them. Another work developed an automatic audit system based on weakly-supervised named entity recognition and deep learning using EMS records and clinical notes, which can potentially reduce the time and labor involved in current manual chart audit reviews [8]. GRACE [5] combined hands-free interfaces with speech recognition and NLP for real-time processing of responders' conversations during the incident and automated filling of ePCR. Another work developed an NLP pipeline to determine treatment appropriateness from EMS motor vehicle crash records [9]. In this study, we conducted a comprehensive analysis of pre-collected ePCR to develop an EMS-specific ontology, analytics, and insights that can facilitate the design of such cognitive assistance systems.

EMS Data Analysis: Previous works on the analysis of EMS data have focused mainly on analysis of the national trends and characteristics of emergency medical services in the U.S. based on NEMSIS data [10]–[12], specific illnesses [13], or types of procedures [14] or incidents. Many works have focused on using EMS data to find patterns and trends to combat the Opioid Epidemic [15] [16] [17]. EMS data was chosen due to its timeliness, geographical indexing, and sizable patient population. Another work applied NLP and machine learning (ML) techniques to predict the Cincinnati Prehospital Stroke Scale (CPSS) for stroke patients based on EMS data [18]. Our study focused on an overall analysis of relationships within different aspects of EMS data rather than focusing on specific incidents.

EMS Data Standardization: Previous works regarding medical data standardization or labeling have focused mainly on establishment of region or state-wide data systems [19] [20] [16], standardization for a public health monitoring strategy [21] [22], or labeling clinically relevant data from EMS reports using NLP methods [23]. The most relevant work is EMSContext [24] that proposed a domain-specific and weakly-supervised method for automated creation and expansion of an ontology of patient's signs and symptoms to be used for automated concept extraction in EMS NLP applications. However, the resulting ontology did not include other types of important EMS concepts such as medications, procedures, impressions, and chief complaints.

EHR Data Analysis: Other relevant works on medical data preprocessing and summarization have focused on medical concept mapping [25], information extraction [26], and section identification [27] in the narrative contents of EHR. However, many of the existing EHR analysis methods are not directly applicable to EMS ePCR due to domain-specific terminology and limited labeled datasets in EMS [24]. Further, the nature, structure, and purpose of these two medical reports are very different. EHR are used to record the patient's medical history over time during in-hospital care and, thus, are more comprehensive and less prone to errors and inaccuracies. On the other hand, ePCR are a one-time report of the patient in emergency distress, collected under pressure during an unfolding incident, and often do not contain confirmed diagnoses or full patient history. Thus, there is a need for developing domain-specific lexicons and analytical methods specific to ePCR.

[1]https://github.com/UVA-DSA/EMS-Pipeline/tree/master/ePCR

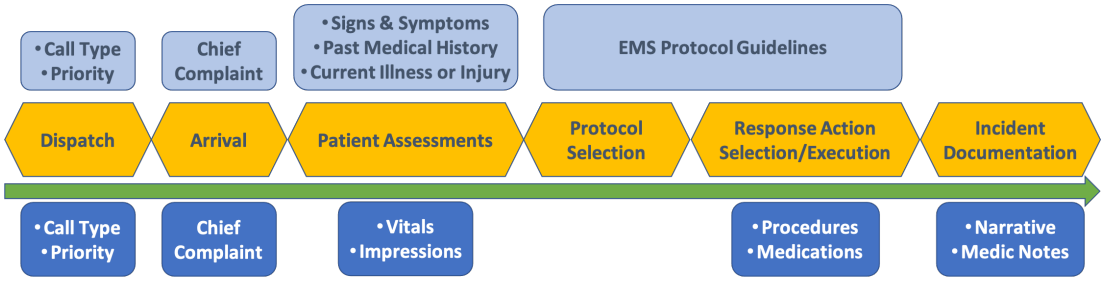

Fig. 1: Overall Progression of an EMS Incident: Steps Taken with Gathered Information and Data Generated for the Reports

III. Background

During every EMS incident, which is usually initiated by a 911 "call," there is a specific set of steps EMS responders follow that produce a flow of information (Figure 1). In each call, responders are dispatched to the incident location and informed of the "Call Type," which is the general reason for the incident. On arrival, responders interact with the patient and others to identify the "Chief Complaints," which are the primary reasons for the EMS call. Then, the responders use the patient's Chief Complaints, Signs and Symptoms, past medical history, history of the present illness (HPI) or injury, and current presentation to form a set of "Impressions." From the Impressions, the responders select and follow the appropriate "EMS Protocol Guidelines" to perform "Interventions" (including "Procedures" and "Medications"), which are a series of treatments to stabilize the patient before transporting to the hospital. Responders document this information flow from Call Type to Chief Complaints to Impressions and, finally, to Interventions, along with other information described as "Narrative" and/or "Medic Notes", in the EMS incident reports. When a report is documented electronically, it is referred to as an electronic Patient Care Report (ePCR).

We collected 35,926 ePCR, including 28,124 reports from the years 2019-2020 and 7,802 from 2017-2018, from a regional ambulance agency. These reports provide valuable insights into the events that occurred during each EMS call. With this vast amount of data, many co-occurrences and relationships between different aspects of the call, including the impressions made and the steps taken by the responders, can be observed to inform the design of future decision support systems and assistive technologies for first responders. In particular, cognitive assistant systems that can automatically collect data and make sense of the observations made at the scene can help responders to correctly understand and process information and reach correct assessments for the most effective response actions [6], [7].

A. Structure of Data and Terminology

As seen in Figure 2, each report in the dataset contains a set of columns describing different aspects of an incident as reported by the responders. The columns include *Priority, Call Type, Chief Complaints, Impressions, Vitals, Medic Notes, Narrative, Procedures*, *Medications*, *Interventions* (representing both *Procedures* and *Medications* in one group), and *Outcome*. All the columns are described as categorical values in textual format. The columns *Chief Complaints*, *Impressions*, *Vitals*, *Procedures*, and *Medications* are a list of these values, separated by delimiters.

B. Raw Data Challenges

The ePCR were not formatted consistently across the years: (1) In the 2017-2018 data, the procedures and medications (both subcategories of interventions) were grouped into a single column while the 2019-2020 data kept them separate. (2) The delimiters that separated values in the columns varied by year between quotes, white-space, commas, and braces. This lack of unified delimiters is in particular problematic when isolating individual concepts within each aspect. (3) The ePCR presented variances in the ontology. The 2019-2020 data included medically standard unique identifiers for each value in the impression, medication, and procedure columns, while the 2017-2018 data did not include any. In addition, there were semantic variations in the data across the years for the same concept. For example, a common chief complaint, chest pain, was represented as "pain in chest," "CP," "my chest hurts," etc. This is problematic because each semantic variation appeared as separate individual concepts, which would artificially increase the number of concepts per aspect during data analysis.

IV. Methods

To address the aforementioned challenges, we completed and corrected the ePCR and created a domain-specific ontology of EMS concepts by developing dictionaries for medications, procedures, impressions, chief complaints, call types, and signs and symptoms (see Figure 2). We adopted a semi-automated approach for the development of ontology, consisting of intelligent pattern and similarity/group detection and concept extraction from unstructured text in combination with background EMS knowledge and external queries to standard medical identifier systems. Each dictionary contains a column of domain-specific unified concepts along with corresponding columns for each concept's unique ID and a seed list of relevant keywords or semantic variations used for describing the concept. In the future, machine learning methods in combination with web crawling and API queries could fully automate this process. Specifically, these EMS-specific dictionaries can be further expanded using automated methods such as those proposed in [24] to capture the variations in terminology used in different datasets and across the EMS domain.

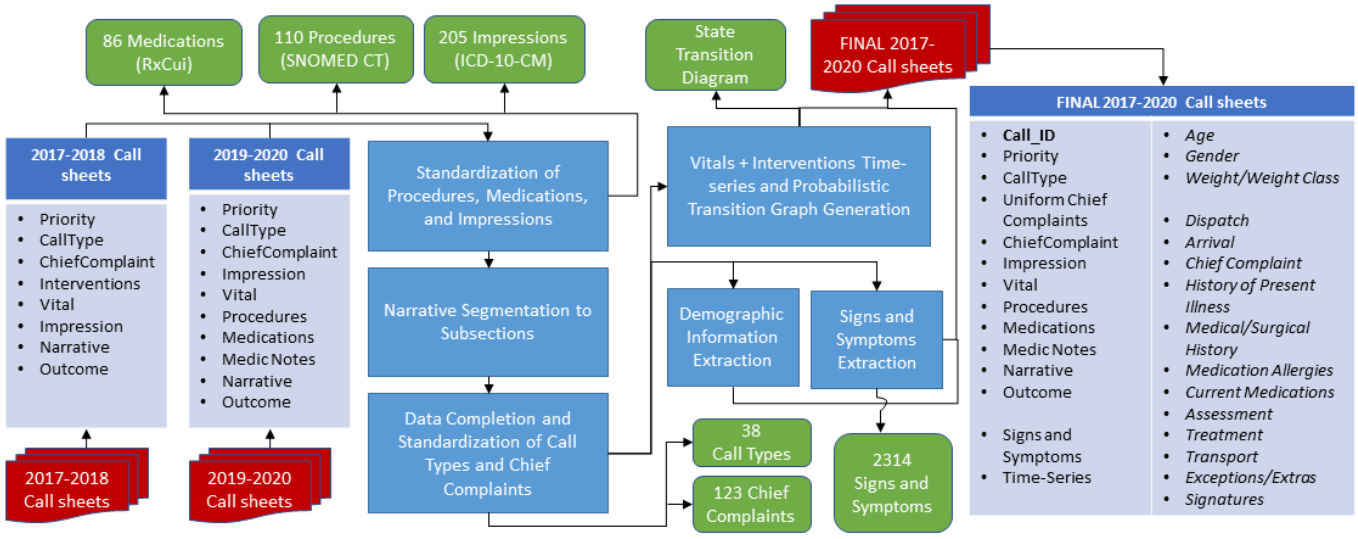

Fig. 2: Overall Method for EMS Data Standardization and Ontology Creation

We then developed methods for automated segmentation of the unstructured text in the *Narrative* and *Medic Notes* of the ePCR and the extraction of additional structured information, including *Chief Complaints*, Signs and Symptoms, and patient demographic information. We also generated time-series representations and state transition diagrams of interventions and vitals. Finally, we conducted an offline analysis of relationships and trends within different aspects of the completed and corrected ePCR. The overall process is summarized in Figure 2 and each step is described in more detail next.

A. Standardization of Procedures, Medications, Impressions

To create a standardized ontology of EMS concepts, we first identified all different textual representations (semantic variations) of the unique concepts mentioned for procedures, medications, and impressions in the reports. We then grouped semantically similar representations of each concept and assigned unique identifiers to each concept group. This was done through rule-based parsing, manual separation and grouping, referencing the regional ambulance agency's medications list, and external queries to publicly-available databases to access the unique identifiers: RxCui [28], [29] for medications, Snomed-CT [30] for procedures, and ICD-10-CM for impressions [31]. These standard identifier systems are widely used within the healthcare domain to categorize and identify medications, procedures, and impressions. They are useful for the EMS ontology as they can provide a common language for clinical terms and seamless integration with other systems such as EHR.

The created dictionaries were then utilized to separate medications from procedures in the 2017-2018 *Interventions* data and to map the various semantic variations to their respective concept in the ontology. The result was one merged dataset of reports from 2017-2020 of unified concepts for medications, procedures, and impressions.

B. Narrative Segmentation to Subsections

The *Narrative* contains the free-form textual description of the observations and actions taken during incidents, documented by the responders after the call. This textual description is organized into several subsections. However, the delimiters for different subsections were not consistent across the dataset and had different forms (e.g., "D-" vs. "D:" vs. "Dispatch:" all represented the Dispatch section). Since the *Narrative* was collected by manually typing into ImageTrend [3] or similar software, an individual responder's personal writing preferences could be one of many possible causes for this disparity. In addition, the total number of subsections and the different aspects of the call they represented were initially unknown to us. We first focused on identifying and defining the different forms of the delimiters. Each form varied in length (abbreviation versus full word), letter case (lower versus upper), and the symbol type and length (none vs. one vs. two or multiple). After careful manual review of the dataset, we determined eight distinct forms of delimiters for subsections as shown in Table I.

To identify the number of subsections and the aspects that they represented, we used regex matching to find any grouping of letters with the same pattern. The seven regular expressions we generated (one for each form except for WLN), are shown in Table I. The regular expressions for the WLN form were generated after determining the aspects represented by each subsection. After filtering out erroneous matches that belonged to other concepts mentioned in the narrative, such as interventions or impressions (e.g. "ECG-" was a match for the format WCO that was erroneous), we were able to find all delimiters for each form. Then, we found the aspect of the call that each

TABLE I: Eights forms of delimiters for the *Narrative*

Form	Abbreviation	Regex expression	Delimiters for the Dispatch subsection, separated by braces
Abbrev lowercase, one symbol	ALO	`(((?<![a-zA-Z0-9\/\.])[a-z]1(?:-|- ))(?![0-9]))`	{d-}
Abbrev capital, one symbol	ACO	`((?<![a-zA-Z0-9\/\.])[A-Z]1(?::|-| -|\/ |\.))`	{D-}{D;}{D.}{D/}{D;}{D)}{D-}{D -}
Abbrev capital, multiple symbols	ACM	`((?<![a-zA-Z0-9\/\.])\([A-Z]1(?:\)-|\)))`	{(D)}{(D)-}
Word capital, one symbol	WCO	`([A-Z]2,[ ]?(?:[A-Z]?\/?)0,(?:-|:))`	{DISPATCH-}{SUBJECTIVE:}
Word capital, two symbols	WCT	`(»[A-Z]2,[ ]?(?:[A-Z]?\/?)0,-)`	{»DISPATCH-}
Word lowercase, one symbol	WLO	`([A-Z][a-z]2,[\/ ]?(?:[A-Za-z]?\/?)0,:)`	{Dispatch:}{Dispatched For:}
Word lowercase, two symbols	WLT	`(\[[A-Za-z \.\/,]+?\])`	{[Dispatch]}
Word lowercase, no symbols	WLN	`(<Aspect>\S*)`	{Dispatched}

subsection represented by mapping the capital or lowercase forms to the most suitable aspect in our EMS ontology (e.g., D- corresponds to "Dispatch:"). Using this method, twelve distinct narrative subsections and corresponding aspects were discovered, including *Dispatch*, *Arrival*, *Chief Complaints*, *History of Present Illness*, *Medical/Surgical History*, *Medication Allergies*, *Current Medications*, *Assessment*, *Treatment*, *Transport*, *Exceptions/Extras*, and *Signatures*. The resulting reference table (partially shown for the *Dispatch* subsection in Table I), containing all different forms of each aspect's delimiters, was used to segment each narrative into its respective aspects.

Segmentation requires consideration of the differences in delimiter formatting and a proper predefined reference table to minimize error. This allows for matching on regex strings that are more literal than pattern-based so that erroneous segmentation is less likely to occur (e.g., matching exactly on "A:" instead of all capital letters preceding a colon). For example, Figure 3 shows a Narrative in ACO form with the delimiters for the *Assesment*, *Treatment*, and *Transport* subsections highlighted in blue. If the delimiters were not found beforehand and the narrative was segmented on any group of characters preceding a punctuation (e.g., colon, dash, slash), there would be many erroneous segmentations as highlighted in gray. If the ACO form of the narrative was not considered, the "d-" shown in red, which is a valid delimiter, but in ALO form, will be incorrectly detected as the delimiter for the *Dispatch* section.

C. Completion and Standardization of Call Types and Chief Complaints

The *Chief Complaints* aspect was missing for 13,552 out of 35,926 (37.7%) reports. However, the Narrative Segmentation process yielded 12 additional aspects for each report, one of which were chief complaints. By extracting the chief complaints from the Narrative segmentation, we were able to complete 10,682 out of 13,552 (79%) missing *Chief Complaints* values. The remaining values could not be completed because the chief complaint was not mentioned in the *Narrative*, or the *Narrative* itself was missing. As a result, only 8% of the reports in the final dataset were missing *Chief Complaints* information. Thus, the Narrative segmentation process could also serve as an online automated method for detection and correction of missing values in the reports.

The completed chief complaints and the call types needed to be standardized due to semantic variations representing the same concepts. The process for standardizing the *Chief Complaints* and *CallType* aspects was similar to the method

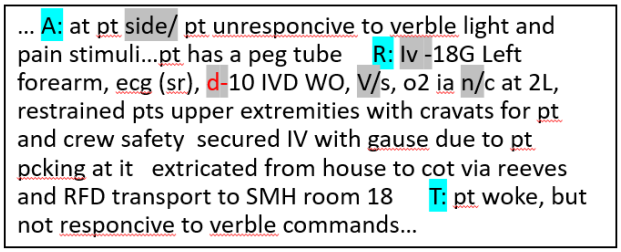

Fig. 3: Examples of Erroneous Delimiter Matches

used for the standardization of procedures, medications, and impressions in Section IV-A).

The call types had a small number of semantic variations for each concept, and less variation within the semantics (e.g., "Fall" vs. "Fall(s)"). This allowed for manual grouping of the semantic variations and using the grouped Call Types dictionary to unify all semantic variations. The chief complaints had more semantic variations for each concept and more variation within the semantics. This is because some chief complaints were recorded verbatim, while others were shortened or simplified by the responders. For example, some chief complaints were recorded as full sentences, e.g., "It feels like someone is stabbing me in my chest and making it hard to breathe," while others were abbreviated (e.g., "sob" for Shortness of Breath). There were a total of 12,079 semantic variations for all chief complaints, so manual review and grouping into unique sets of concepts would have been very time and labor-intensive.

Instead, an alternative method was used. Since the semantic variations for *Chief Complaints* were highly uniform in the reports from 2017-2018, the concepts in the chief complaint dictionary were identified by manually reading the data from only those years. Then, each semantic variation in the 2019-2020 data was automatically mapped using specific "keywords" that were associated with the concept using regex matching. For example, a semantic variation was mapped to "Chest Pain" if the expression contained, in any order, the word "chest" along with any synonyms of pain – throbs, hurts, killing me, burn, pressure, etc. Using this automatic mapping, semantic expressions such as "chest is heavy," "pressure on chest," or "my chest is killing me" were all mapped to "Chest Pain."

D. Demographic Information Extraction and Correction

Demographic information of the patient (i.e., age, gender, weight, and weight class) was extracted from the *Narrative* and the *Medic Notes* when available. Automatically identifying the demographic information relied on mapping "keywords" that

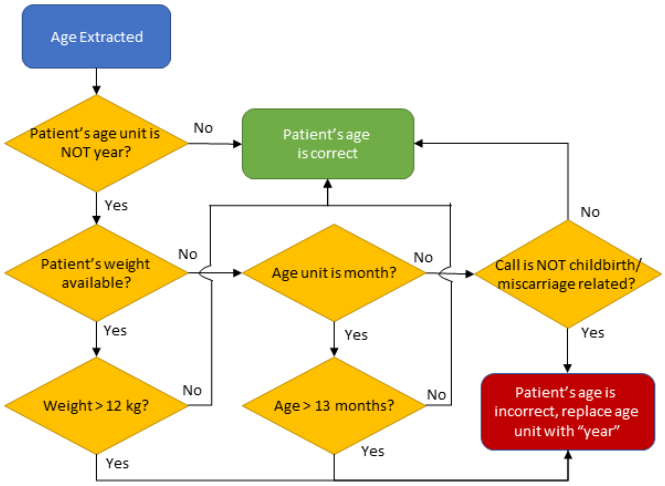

Fig. 4: Correction of Age Units during Age Extraction

related to age, gender, and weight using regex matching. For gender, we searched for matches to variations of gendered pronouns to determine if the patient was male or female (e.g., "she/her" versus "he/his"). For weight, we searched for variations of weight units (e.g., "KG," "lbs," "pounds") and extracted the numbers directly left of the matches. For weight class, we searched for variations and degrees of obesity ("overweight," "morbidly obese") to determine if the patient was obese or morbidly obese. For age, we searched for variations of age units (e.g., "years old," "y/o," "minutes"), then extracted the values associated with the units and corrected any errors in the extracted age information. The correction of demographic information in the reports is particularly important because different age and gender groups require the execution of different treatment protocols and interventions (e.g., different medication dosages).

We observed that the *Narrative* sometimes incorrectly reported the patient's age to be much younger by recording an incorrect age unit (i.e., hour, week, or month instead of year), while from the context of the narrative it was obvious that the patient was an adult. For example, one of the reports stated that the patient was 77 days old and was "talking to RFD and family," which is very unlikely for a 2 month-year-old to do. As a solution, the automated process seen in Figure 4 was developed to detect and correct a total of 176 reports that had incorrect age units. The process consisted of verifying if the age unit reported fits in the context of the entire narrative. First, if the patient's age was not in years, there was a possibility that the age unit was incorrect. Three factors make an age unit smaller than "year" fit in the narrative context: (1) If the weight of the patient is greater than 12 kg, then the patient is most likely an adult. Patients less than 1-year-old typically weigh less than 12 kg. (2) If the patient's age is in a unit smaller than "months," but there is no reference to "Pregnancy," "Childbirth," or "Miscarriage" in the Narrative, the patient is most likely an adult. When the patient is less than 1 month old, the patient has typically been newly born or spontaneously aborted. (3) If the patient's age unit is in months, but the numerical value of the age is greater than 23, the patient is most likely an adult. After a young age of 2 years or 24 months, age is typically expressed in the largest unit possible (e.g., it is more common to express an age as "4 years" instead of "48 months.")

E. Signs and Symptoms Extraction

Signs and Symptoms are the key information collected by the responders in forming impressions and selecting appropriate protocols and response actions. This information is available as part of the *Narrative* and *Medic Notes* in unstructured text format. We adopted a widely-used biomedical text annotation tool called MetaMap [25] from the National Library of Medicine [32], combined with a domain-specific ontology of EMS concepts [6], [24] and semantic type filtering to extract medical concepts corresponding to Signs and Symptoms from the narratives.

Metamap uses symbolic and computational linguistic techniques and natural language processing for mapping the biomedical text into concepts in the Unified Medical Language System (UMLS) Metathesaurus. UMLS is the largest thesaurus in the biomedical domain which provides a representation of biomedical knowledge using concepts with different semantic types and relationships [25]. Every UMLS concept is assigned with a unique identifier called Clinical Unique Identifier (CUI). Using Metamap we extracted all the biomedical concepts in the narratives along with their semantic types, CUIs, negation condition, and position in the text. There are about 126 different kinds of semantic types, but we only considered the semantic types related to "signs and symptoms" and filtered out the rest. We also filtered those concepts that were indicated as negated by Metamap. After filtering, we ended up with a dictionary of over 2,000 unique signs and symptoms as shown in Figure 2. We further merged similar signs and symptoms in this dictionary into unique concept classes based on semantic similarity and using the EMS ontology from [6], [24].

F. Treatment Sequence Analysis

Responders follow established guidelines for assessing patients' conditions, identifying the most appropriate treatment protocols, and then executing different intervention steps according to the protocols. However, identifying the set of protocols considered and executed by the responders based on ePCR information is challenging [6]. Tracking the call progression over time through analysis of time-series data can provide important insights about the sequence of observations made by the responder and the corresponding intervention steps. We processed the *Vitals, Procedures,* and *Medications* data and created a unified and chronologically sorted time-series representation of the important observations and interventions during the call. This time-series is represented as a list of tuples, each formatted as:

(Aspect, Unified Concept, Semantic Variation, Value, Time)

representing the type of aspect (i.e., signs and symptom, vital sign, procedure, or medication), its value represented

(Medication, Oxygen, Oxygen, N/A, 07:31:06),
(Procedure, Cardiac Monitor, CV - ECG - 4 Lead, N/A, 07:37:45),
(Vital, Pulse, Pulse, 108, 07:37:45), ...
(Vital, Electrocardiogram, EKG, N/A, 07:37:45),
(Procedure, IV, IV Start - Extremity Vein, N/A, 07:43:15),
(Medication, Naloxone, Naloxone (Narcan), N/A, 07:53:01),
(Vital, Pulse, Pulse, 94, 07:58:12), ...
(Vital, Electrocardiogram, EKG, Sinus Tachycardia, 07:58:12)

Fig. 5: Example EMS Incident Time-series

as a unified concept with a unique identifier, the semantic variation used in the narrative, the numerical value (for the vital sign values or medication dosages), and the time that the observation was made or the action was taken.

For example, we can observe in the time-series of Figure 5 that the responder decided to administer oxygen after the initial encounter, even before taking vitals. This indicates that there was a clear sign of the patient being in respiratory distress or hypoxic. Then, the responder took standard vitals and an ECG. From that information, the responder was able to form an impression and select the appropriate protocol. The patient's first set of vitals and the interventions taken indicate that the responder followed the protocol for a Drug Overdose. After completing the protocol, the responder took a second set of vitals to reassess the patient's status.

By analyzing the time-series data across all the calls, we can see if there is a common sequence of intervention paths taken by responders for specific patient conditions. This analysis combined with the models of protocols can identify the similarities in the pathways taken by the responders, and any possible deviations from the protocols. This information will be helpful for both cognitive assistance and performance monitoring [6].

After converting the ePCR into a unified time-series format, we created a state transition graph representing the sequences and probabilities of intervention steps taken for each impression. To find the state transition graph, each sequence of interventions treating an impression must be combined into one sequence with transition probabilities. Each sequence, represented as a time-series, can be treated as an unweighted directed graph with each vertex representing an extracted concept and the edges representing the ordering of concepts in the time-series. To extract a state transition diagram based on all the calls, we found the union of all the unweighted directed graphs generated for a specific impression. To do so, the weight of each edge was increased by one for each repetition of an edge in a directed graph. After the union was found, the weights were converted into probabilities by summing the weights of all edges going out of a vertex, then dividing the weight of each edge going out of the same vertex by the sum. The result was a weighted directed graph, with the weights representing the transition probabilities or the percentage of time that a specific sequence of actions was performed by the responders. The pseudocode that further specifies the process is shown in Algorithm 1. The weighted directed graph G is represented as a hashmap of hashmaps. Each vertex v in the graph is stored in the outer hashmap's key and the set of the vertex's adjacent vertices U and their corresponding $weights$ are stored in an inner hashmap representing the adj_list of that vertex.

Algorithm 1 Converting Time-series to Adjacency List

```
Impressions ← Set of all unique Impressions in the dataset
for im ∈ Impressions do
    G ← {}
    Calls ← Set of all the rows in ePCR that contain Impression im
    for c ∈ Calls do
        Int ← Ordered list of Interventions for one row from ePCR c
        for i ← 0 to len(Int) − 2 do
            v ← Unique ID for Int[i]
            if v ∉ G then
                G[v] ← {}
            u ← Int[i + 1]
            if u ∉ G[v] then
                G[v][u] ← 1
            else
                G[v][u] += 1
    for v ∈ G do
        U ← v.adj_list
        s ← 0
        for u ∈ U do
            s+ = weight(u)
        for u ∈ U do
            adj_list[u] ← weight(u)/s
```

V. Results

We used the corrected, completed, and expanded structured information from ePCR to perform further analysis of the trends in ePCR over the years. In our analysis, we determined the relationship among different aspects of incident response by calculating the conditional probability *P(X|Y)* of one aspect given another. Using conditional probability allowed us to infer the probability of responders conducting aspect *X* when they know or conclude on aspect *Y*, as indicated by *P(X|Y)*. For instance, we can infer with *P(Procedure|Impression)* the probability that a responder will conduct a specific procedure when they make a certain impression about a patient. Each cell *(X, Y)* in the heatmaps shown in this section corresponds to the value of *P(X|Y)*. In addition, we determined the treatment sequence for impressions by generating state transition maps with the ePCR as described in the Methods section. In most cases, our analyses agreed with knowledge well-known to responders, which shows that corrected and completed ePCR could be of use for data-driven approaches. The main results and observations from this analysis are summarized next.

A. Co-occurrence of Impressions

Since impressions sometimes occur in multi-value associated pairs in one single call, we determined the conditional probability of one impression given one out of the top 20 frequently occurring impressions *P(Impression|Impression)* to identify trends in the co-occurrence of impressions. Figure 6 shows the major co-occurrence frequencies (greater than or equal to 10%) for the top 20 frequent impressions, which account for 61% of all occurrences of impressions in the data.

Impressions come in strongly occurring pairs. For instance, Altered Mental Status was strongly associated (co-

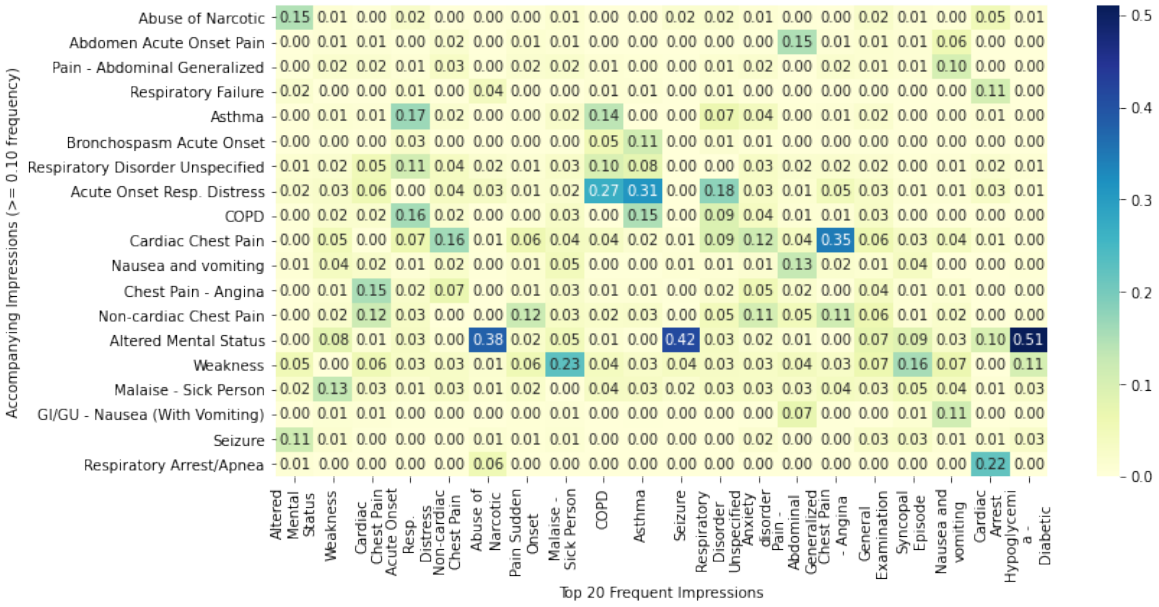

Fig. 6: Co-occurrence of Different Impressions

occurring 51% of the time) with the occurrence of Diabetic Hypoglycemia. Co-occurrence of several types of impressions shows the challenge of the differential diagnosis process conducted by responders. For example, seizures are typically present with altered mental status, but altered mental status also frequently occurs in situations not involving seizures.

B. Procedures vs. Impressions

After grouping similar procedures into groups (e.g., grouping "BVM via mask" and "BVM via tube" as "BVM"), we determined the conditional probability of a procedure given a top 20 frequent impression *P(Procedure|Impression)*. Procedures with significant co-occurrence frequencies (greater than or equal to 5%) are shown in Figure 7. Irrelevant procedures, or most common procedures conducted in almost every call with different impressions (e.g., ECG, IV/IO, Vitals/Asses, Monitor/Care, ETCO2 Capnography, and Move Patient), were not included in this analysis because they did not provide distinguishing information between the impressions.

Frequencies of procedures were different based on the impression. Certain procedures were conducted more frequently in some impressions than others. For instance, Nebulizer Therapy was conducted very frequently at 51% and 63% for Impressions COPD and Asthma, while conducted with less than 1% frequency in other impressions such as Syncopal Episodes, Abuse of narcotic, and Pain Sudden Onset. This is consistent with the knowledge of EMS protocol guidelines that indicate Nebulizer Therapy as the first-line treatment for COPD and Asthma. This indicates that the procedures conducted, and therefore which protocol to follow, were decided based on the responder's impression of the patient's condition. Therefore, it is critical for assistive technologies to infer the impressions made by the responders and accurately model the decision flow from impressions to procedures.

C. Medications vs. Impressions

We determined the conditional probability of a medication administration given a top 20 frequent impression *P(Medication|Impression)*. Significant co-occurrence frequencies (of greater than or equal to 5%) for medications and impressions are shown in Figure 8.

We observed that certain medications were administered more frequently in some impressions than others. For instance, the medication Epinephrine 1:10,000 had a 62% chance of co-occurrence in Cardiac Arrest cases but had a frequency of no more than 7% for other impressions. This indicates that medications administered, and therefore, which protocols to follow, were decided based on the responder's impression of the patient's condition and past medical history. Therefore, it is critical for assistive technologies to arrive at correct impressions and accurately model this decision flow from impressions to medications.

D. Chief Complaints vs. Call Types

There were 38 distinct *Call Types* and 123 *Chief Complaints* in the dataset. After grouping similar Chief Complaints based on similarity in conditions (e.g., grouping "Cardiac palpitations," "Congestive Heart Failure," and "Chest Pain" as "Cardiac Pain/Problems"), we determined the conditional probability of a chief complaint group given a call type *P(Chief Complaint|Call Type)*. Only significant frequencies of greater than 15% were considered.

We observed a general correspondence between *Call Type* and *Chief Complaints* values. However, more generalized Call Types diverged into more specific chief complaints. For instance, while "Sick Person" most frequently corresponded with the "General/Other Illness" chief complaints, it also diverged into four different Chief Complaints (Respiratory/Pulmonary Problem, Altered Mental State, and Pain) with a frequency

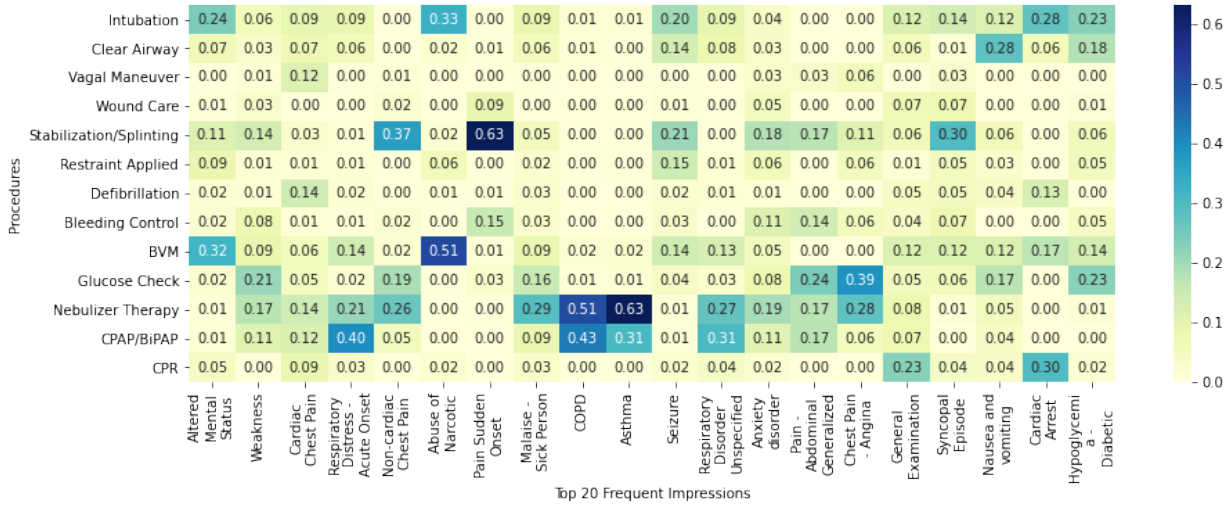

Fig. 7: Co-occurrence of Grouped Procedures and Impressions

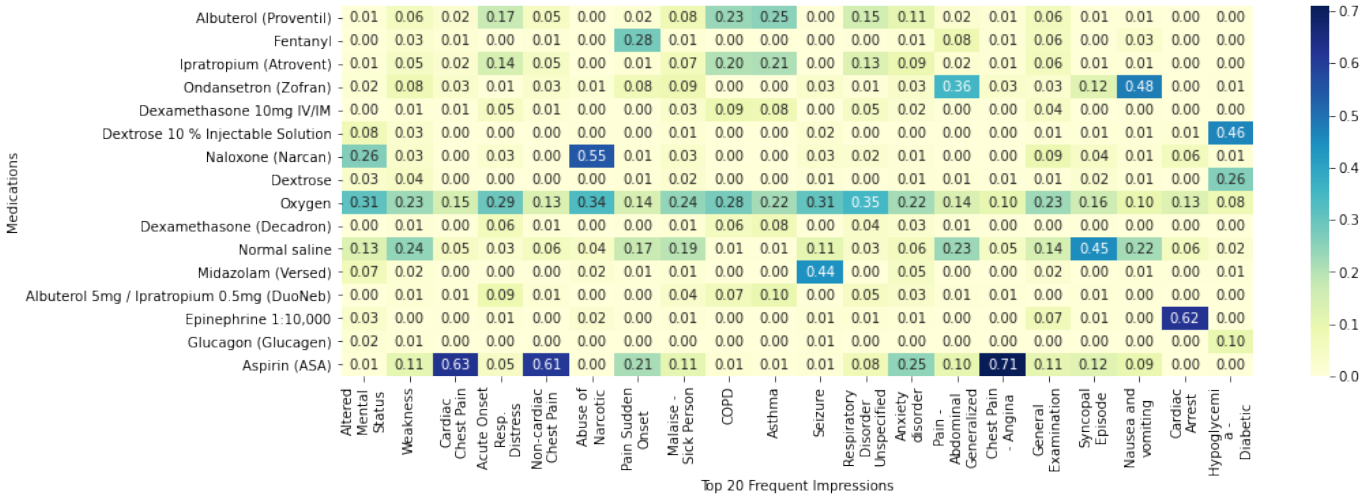

Fig. 8: Co-occurrence of Medications and Impressions

greater than 10%. This is expected as the "Sick Person" calls only suggest a medical problem when the 911 call taker is unable to extract any better information from the caller. Face-to-face responder assessment of the patient is then needed to collect signs and symptoms to determine the chief complaint. In addition, cause-related call types such as Electrocution/-Lightning, Water Accident, and Traffic/Transportation Incident diverged more frequently into medical chief complaints rather than traumatic or environmental chief complaints. Call Types may give responders a general idea of the incident before arrival. However, the Call Type is not a guarantee of the patient's situation due to its generalized nature. Thus, an EMS assistive technology must be able to hone in on the specific issue based on surrounding information, and not be biased by eliminating certain pathways based solely on the Call Type.

E. Grouped Impressions vs. Grouped Chief Complaints

Our analysis identified a total of 123 unique Chief Complaints and 205 Impressions. We grouped the concepts in each aspect based on semantic and condition similarity (e.g., grouping the impressions such as "Childbirth Uncomplicated" and "Preterm Labor with Preterm Deliver" as "OBGYN/Pregnancy/Birth"), and then determined the conditional probability of an impression group given a chief complaint group *P(Impression|Chief Complaint)*.

We observed a general correspondence between Chief Complaints and Impressions. However, more generalized Chief Complaints often diverged into more specific Impressions. For instance, the very general group of chief complaints "Other" corresponded to the potentially life-threatening impression "Cardiac Pain/Problems" at a frequency of 10%. This indicates that there may be underlying causes to the problem that the patient is not aware of and/or does not state. Although the Chief Complaint is important for identifying the patient's situation, it should not be the only information considered. To be most efficient, EMS assistive technology must consider life-threatening conditions, even when not reported by the patient, based on other information such as signs and symptoms,

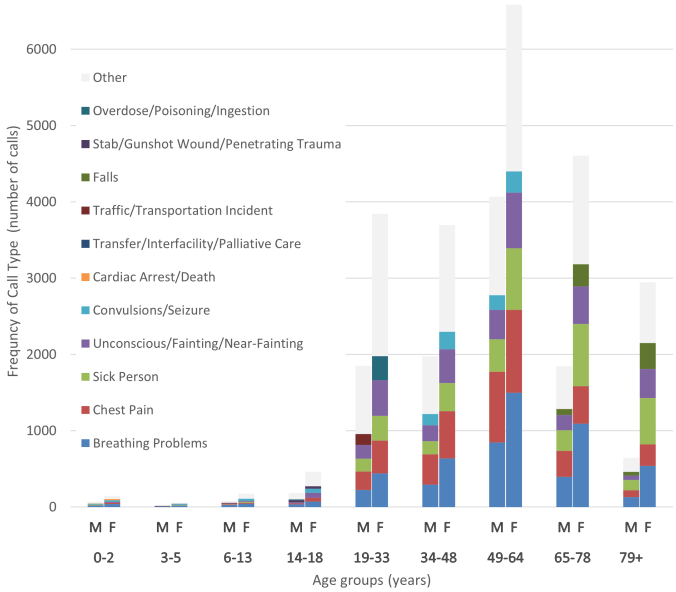

Fig. 9: Call Type Frequency vs. Age Group and Gender

bystander reports, and not solely based on the patient's Chief Complaint.

F. Call Type Frequency vs. Demographics

Using the extracted demographic information, we were able to investigate the relationship between age group gender and Call Type. Figure 9 shows the top 5 frequent call types that occurred for each age group and gender. The "Other" category includes all the call types that occurred less frequently than the top 5 categories in each demographic group. These call types each were associated with a very small number of ePCR and did not hold enough significance to be considered individually in our analysis, so we grouped them together.

The age group and gender with the greatest number of calls were 49 to 64-year-old females. We consistently observed 2-3 times more female than male patients across all age groups. The difference was especially high for ages 79+. This is consistent with the common knowledge of responders that women are more likely to call for assistance with medical problems than men. In addition, women typically live longer than men, and this was reflected in the data by the increasing difference in calls for men versus women as the age group increased. For 19 to 33-year-old female patients, calls for "Overdose/Poisoning" were frequent, while for males in the same age group "Traffic/Transportation incidents" were frequent. This is consistent with poisoning being the preferred mechanism for suicide among females, and trauma being the preferred mechanism for males. Also, males tend to drive more and take more chances than females. In addition, the "Stab/Gunshot Wound/Penetrating Trauma" Call Type was included in the top 5 for only females and males of ages 14-18. This observation is likely related to gang activity as the dataset is from an urban agency. Also, this age group typically experiences more traumatic injuries than medical problems in most settings, exacerbated by an urban setting with higher levels of conflict, crime, and risk-taking. "Convulsions/Seizures" were frequent in ages 34-64 for both genders. Although some patients have life-long seizure conditions with known causes and treatment, a new-onset seizure among adults is taken seriously and typically results in a 911 call. A seizure can be scary for family, friends, and bystanders, and the first instinct is to call for help. "Falls" were the most frequent for ages 65+ as aging can reduce balance, flexibility, and reaction time and older people are more likely to fall. Physical resilience also decreases with age, so falls are more likely to cause injury.

While age and gender are not definite indications of the patient's situation, demographics are an important factor for responders to accurately assess a patient's conditions and needs. Demographics is an important part of the patient profile and should be included in models of assistive technologies to increase their effectiveness.

TABLE II: Frequent Impressions and top related Signs/Symptoms

Impressions mapped to Protocols	**Top 5 Related Signs and Symptoms**
Asthma/COPD/Croup	respiratory distress, wheezing, coughing, hypertensive disease, sinus tachycardia
Fever	respiratory distress, weakness,sinus tachycardia, coughing, vomiting
Hypoglycemia	hypoglycemia, diabetes, hypertensive disease, confusion, ablepharonmacrostomia syndrome
Seizures	tremor,confusion, ablepharonmacrostomia syndrome, sinus tachycardia, cerebrovascular accident
Stroke	cerebrovascular accident, facial paresis, weakness, hypertensive disease, headache
Allergic Reaction	pruritus, respiratory distress, urticaria, welts, ataxia telangiectasia
Abdominal Pain	vomiting, nausea, chest pain, diarrhea, respiratory distress
Altered Mental Status	ablepharonmacrostomia syndrome, confusion,seizure, sinus tachycardia
Sepsis	weakness, sinus tachycardia, respiratory distress, confusion, ablepharonmacrostomia syndrome
Congestive Heart Failure	respiratory distress, rales, hypertensive disease, COPD, chest pain

G. Signs and Symptoms vs. Impressions

We determined the conditional probability of impressions given signs and symptoms to understand better the key signs and symptoms that affect responders' impressions and choice of protocols as seen in Table II. We firstly choose 10 groups of impressions that could be easily mapped to specific EMS protocols. For example, the impression "hypoglycemia" can be directly mapped to the protocol "Medical – Diabetic – Hypoglycemia". There were two reasons for mapping impressions to protocols: (1) It would be possible to verify the relationships between signs and symptoms and impressions by looking up the mapped protocol's description from EMS protocol guideline documents. (2) Responders are expected and are trained to make decisions based on the established protocol guidelines. Therefore, in this section we will refer to these 10 groups of impressions as protocols they are mapped to. In Table II, the top 5 related signs and symptoms of each protocol are listed, and they are ranked by the conditional probability values *P(Sign and Symptom|Impression)*.

Our analysis showed that: (1) Some concepts extracted as signs and symptoms from the narratives exactly matched with the impressions documented by the responders and names of specific protocols (e.g., seizure, fever, sepsis, and cardiac arrest). This shows that more specific semantic filtering of extracted concepts might be needed. (2) Most of the top 5

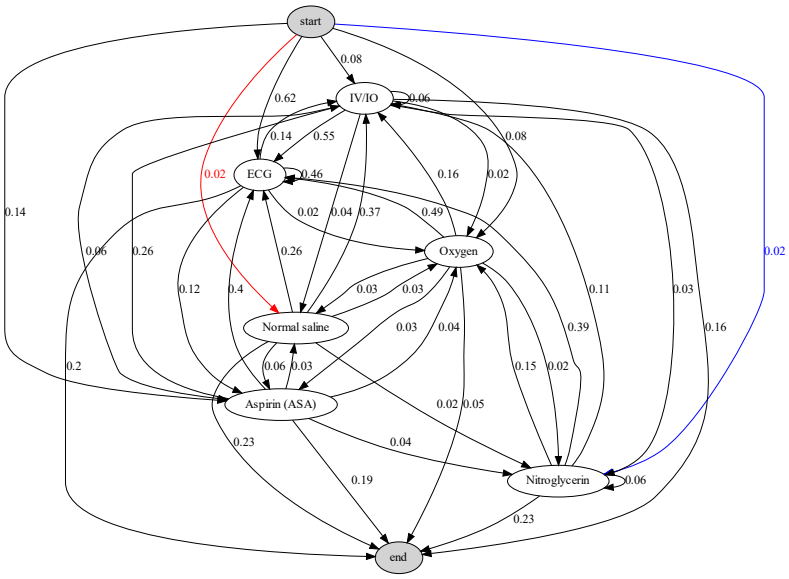

(a)

1. Perform general patient management.
2. Support life-threatening problems associated with airway, breathing, and circulation.
3. Administer oxygen to maintain *SPO₂* 94 - 99%
4. Establish an IV of normal saline per patient assessment.
5. Obtain 12 lead ECG.
   a. If 12 lead reads, "***AMI***, the patient should be immediately transported to the closest PCI capable hospital. AIC must notify receiving facility ASAP.
   b. If 12 lead is consistent with STEMI, and capability exists, transmit 12 lead to PCI center.
6. Transport immediately.
   a. If actual transport time is greater than 45 minutes to a PCI center, consider use of aeromedical.
7. Place patient on cardiac monitor and monitor pulse oximetry.
8. If no contraindications, administer ASA 324 mg PO.
9. If confirmed STEMI and/or significant cardiac history, administer ***NITROGLYCERIN*** 0.4 mg SL. If the pain persists and B/P > 100 mmHg systolic, repeat nitroglycerin 0.4 mg SL in 3 to 5 minutes (up to total of three SL doses).
10. If pain persists, refer to *General – Pain Control* protocol.
11. Transport and perform ongoing assessment as indicated.

(b)

Fig. 10: (a) Sequence and Frequency of Treatment Steps for STEMI/MI Chest Pain Protocol, (b) ST Elevation Myocardial Infarction (STEMI) Protocol, adopted from Old Dominion EMS Alliance (ODEMSA)

related signs and symptoms found for the protocols agreed with the indications of protocols provided in the protocol guideline documents. For example, the most related sign and symptom for protocol "Asthma/COPD/Croup" was "respiratory distress", which means if the responder reported an impression of "Asthma/COPD/Croup", the patient was most likely observed in respiratory distress. However, there were also some discrepancies between the data and the guidelines. For example, "Ablepharonmacrostomia Syndrome" was a sign and symptom that is not listed in the indications of some protocols such as "Seizure", however, our data analysis showed this syndrome among the top 3 related signs and symptoms for the "Seizure" protocol. (3) "respiratory distress" is the top 1 related sign and symptom for both protocols "Asthma/COPD/Croup" and "Cognitive Heart Failure", which means responders need to rely on other signs and symptoms to distinguish between these two protocols. These insights are helpful for the design of automated protocol prediction algorithms that can assist responders in decision making [6]. Specifically, the discrepancies between data and knowledge sources show the possibility of errors and rare scenarios in the data and that we need to rely on combined model and data-driven methods for accurate protocol identification.

H. Treatment Sequence Analysis for Different Impressions

We created state transition graphs representing the frequencies of different treatment sequences taken for each group of impressions that could easily be mapped to specific EMS protocols. Figure 10a shows an example transition graph with the different treatment paths that responders took for the impressions relating to STEMI/MI Chest Pain protocol, with the paths containing treatment steps greater than a probability of 2% for clarity. Figure 10b shows the description of the same protocol based on the EMS guideline documents.

We observed that most treatment sequences in the transition graphs of different protocols contained some form of ECG and IV procedures followed by 1 to 3 medication administrations. In general, the treatment steps taken in different calls mostly followed the protocol guidelines, as expected. However, there also existed discrepancies between the learned transition graph models and the guideline protocols. For instance, the major paths STEMI/MI Chest Pain protocol in Figure 10a mostly follow the guideline protocols in 10b. However, the treatment step (in red) that goes directly from "start" to "Normal saline," bypassing "IV/IO," deviates from protocol. Normal saline is administered through an IV connection. Since this is common medical knowledge, it is highly likely that in some reports the "IV/IO" procedure is assumed, leading to this artificial path. In these cases, paths that do not agree with common medical knowledge could be used to automatically fill in these assumptions, correcting and completing the data. The treatment step (in blue) that goes directly from "start" to "Nitroglycerin" also deviates from protocol since aspirin should be administered before nitroglycerin. Possibly, aspirin administration may have been not recorded in the report. However, the responder may also have had a justification to "deviate" due to a contraindication, e.g., the patient was allergic to aspirin, or the patient had already taken an aspirin. In these cases, the protocol guideline specifies against aspirin administration. So given the context of the incident, the "deviation" is justified and follows the protocol.

The analysis of consistencies and discrepancies of the protocol transition graphs versus guideline documents could be utilized for: (1) automatic generation of behavioral models

of protocol from data that will help encode the knowledge of protocols in the design of EMS assistive technologies [6], (2) detection of incomplete or incorrect information in the reports, and (3) runtime verification of protocol execution and performance monitoring during EMS training. In addition, it is important to consider the different branches with exceedingly small transition probabilities (<2%) such as those omitted from Figure10b for clarity. Even though numerically these transitions may not seem significant, they may represent the special cases of medical significance that should be considered and learned from. Finally, the many different paths and sequences indicate that the responders and any supporting cognitive assistant technologies must be adaptable to exceptions and rare scenarios that might influence the protocol selection.

VI. Conclusion

This paper addressed challenges in summarizing the substantial amount of data reported on EMS incidents by developing methods for translating reported information into actionable knowledge for performance monitoring and improved operations. We semi-automatically completed, corrected, and then analyzed a large dataset of pre-collected pre-hospital electronic Patient Care Reports (ePCR) from an ambulance agency and developed a domain-specific ontology of EMS concepts. Our analysis provided insights on the relationships among different incident aspects, including patients' chief complaints, signs and symptoms, responders' impressions, and interventions, and evaluated the most common response action sequences. These insights can aid in the design of future EMS assistive technologies to provide decision support and reduce cognitive load for responders. The EMS-specific ontology and analytic methods resulted from this study can be used for both offline and online summarization of ePCR, automated detection and correction of missing or incorrect information in the reports, and generation of models of EMS protocols.

References

[1] "Review of nationwide 911 data collection," July 2013. [Online]. Available: https://www.911.gov/pdf/National_911_Program_Review_Nationwide_Data_Collection_2013.pdf

[2] D. W. Spaite, T. Hanlon, E. A. Criss *et al.*, "Prehospital data entry compliance by paramedics after institution of a comprehensive ems data collection tool," *Annals of emergency medicine*, vol. 19, no. 11, pp. 1270–1273, 1990.

[3] "ImageTrend," https://www.imagetrend.com/.

[4] "Emergency Department Information Exchange," https://collectivemedical.com/impact/ed-utilization/.

[5] "Grace: Generating summary reports automatically for cognitive assistance in emergency response," vol. 34. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/7049

[6] S. Shu, S. Preum, H. M. Pitchford *et al.*, "A behavior tree cognitive assistant system for emergency medical services," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6188–6195.

[7] S. Preum, S. Shu, M. Hotaki *et al.*, "CognitiveEMS: A Cognitive Assistant System for Emergency Medical Services," in *Special Issue on Medical Cyber-Physical Systems Workshop (CPS-Week 2018)*, vol. 16, no. 2. ACM SIGBED Review, 2019.

[8] W. Han, W. Yeung, A. Tung *et al.*, "An emergency medical services clinical audit system driven by named entity recognition from deep learning," 2020.

[9] C. J. Tignanelli, G. M. Silverman, E. A. Lindemann *et al.*, "Natural language processing of prehospital emergency medical services trauma records allows for automated characterization of treatment appropriateness," *The journal of trauma and acute care surgery*, vol. 88, no. 5, p. 607—614, May 2020. [Online]. Available: https://doi.org/10.1097/TA.0000000000002598

[10] H. V. Duong, L. N. Herrera, J. X. Moore *et al.*, "National characteristics of emergency medical services responses for older adults in the united states," *Prehospital emergency care*, vol. 22, no. 1, pp. 7–14, 2018.

[11] H. E. Wang, N. C. Mann, K. E. Jacobson *et al.*, "National characteristics of emergency medical services responses in the united states," *Prehospital emergency care*, vol. 17, no. 1, pp. 8–14, 2013.

[12] J. N. Carlson, C. Karns, N. C. Mann *et al.*, "Procedures performed by emergency medical services in the united states," *Prehospital Emergency Care*, vol. 20, no. 1, pp. 15–21, 2016.

[13] J. Friedman, L. Beletsky, and D. L. Schriger, "Overdose-related cardiac arrests observed by emergency medical services during the us covid-19 epidemic," *JAMA psychiatry*, vol. 78, no. 5, pp. 562–564, 2021.

[14] H. E. Wang, N. C. Mann, G. Mears *et al.*, "Out-of-hospital airway management in the united states," *Resuscitation*, vol. 82, no. 4, pp. 378–385, 2011.

[15] L. Lasher, A.-C. Jason Rhodes MPA, and S. Viner-Brown, "Identification and description of non-fatal opioid overdoses using rhode island ems data, 2016–2018," *Rhode Island Medical Journal*, vol. 102, no. 2, pp. 41–45, 2019.

[16] J. Pesarsick, M. Gwilliam, O. Adeniran *et al.*, "Identifying high-risk areas for nonfatal opioid overdose: a spatial case-control study using ems run data," *Annals of epidemiology*, vol. 36, pp. 20–25, 2019.

[17] A. Garza and S. Dyer, "Ems data can help stop the opioid epidemic," *Journal of Emergency Medical Services*, vol. 41, no. 11, 2016.

[18] R. Garg, C. T. Richards, A. Naidech *et al.*, "Abstract tp296: Predicting cincinnati prehospital stroke scale components in emergency medical services patient care reports using natural language processing and machine learning," *Stroke*, vol. 50, no. Suppl_1, pp. ATP296–ATP296, 2019.

[19] G. D. Mears, D. Pratt, S. W. Glickman *et al.*, "The north carolina ems data system: a comprehensive integrated emergency medical services quality improvement program," *Prehospital Emergency Care*, vol. 14, no. 1, pp. 85–94, 2010.

[20] S. M. Joyce and D. E. Brown, "An optically scanned ems reporting form and analysis system for statewide use: Development and five years' experience," *Annals of emergency medicine*, vol. 20, no. 12, pp. 1325–1330, 1991.

[21] M. M. Rathore, A. Ahmad, A. Paul *et al.*, "Real-time medical emergency response system: exploiting iot and big data for public health," *Journal of medical systems*, vol. 40, no. 12, pp. 1–10, 2016.

[22] T. Krafft, L. G. Castrillo-Riesgo, S. Edwards *et al.*, "European emergency data project (eed project) ems data-based health surveillance system," *The European Journal of Public Health*, vol. 13, no. suppl_1, pp. 85–90, 2003.

[23] G. Silverman, E. Lindemann, G. Rajamani *et al.*, "Named entity recognition in prehospital trauma care," in *MEDINFO 2019*, ser. Studies in Health Technology and Informatics, B. Seroussi, L. Ohno-Machado, L. Ohno-Machado *et al.*, Eds. IOS Press, Aug. 2019, pp. 1586–1587.

[24] S. M. Preum, S. Shu, H. Alemzadeh *et al.*, "Emscontext: Ems protocol-driven concept extraction for cognitive assistance in emergency response." in *AAAI*, 2020, pp. 13 350–13 355.

[25] A. R. Aronson, "Effective mapping of biomedical text to the umls metathesaurus: the metamap program." in *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 2001, p. 17.

[26] M. Devarakonda and C.-H. Tsou, "Automated problem list generation from electronic medical records in ibm watson," in *Twenty-Seventh IAAI Conference*, 2015.

[27] A. Pomares-Quimbaya, M. Kreuzthaler, and S. Schulz, "Current approaches to identify sections within clinical narratives from electronic health records: a systematic review," *BMC Medical Research Methodology*, vol. 19, no. 1, p. 155, Jul 2019. [Online]. Available: https://doi.org/10.1186/s12874-019-0792-y

[28] "RxNORM," https://bioportal.bioontology.org/ontologies/RXNORM.

[29] "RxNav," https://mor.nlm.nih.gov/RxNav/.

[30] "SNOMED CT," https://bioportal.bioontology.org/ontologies/SNOMEDCT.

[31] "ICD-10-CM," https://www.icd10data.com/.

[32] "National Library of Medicine," https://www.nlm.nih.gov/.