PATHWAY EXPLORATION FOR BREAST CANCER CARE

By

Ravi Viswanathan Atreya

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Biomedical Informatics

May 31, 2017

Nashville, Tennessee

Approved:

Mia A. Levy, M.D., Ph.D.

Thomas A. Lasko, M.D., Ph.D.

Joshua C. Denny, M.D., M.S.

Daniel Fabbri, Ph.D.

R. Daniel Beauchamp, M.D.

ACKNOWLEDGEMENTS

I would like to thank my dissertation advisor, Dr. Mia Levy, for her mentorship and guidance through this journey. She has provided great feedback and helped me think about the role of informatics in oncology in new ways. I also appreciate her advice on developing my career. I would also like to thank my committee members Drs. Tom Lasko, Josh Denny, Daniel Fabbri, and Daniel Beauchamp. I enjoyed discussing how best to let the data speak via various data mining and machine learning methods with Dr. Lasko. I appreciated Dr. Denny sharing his insights on phenotyping and natural language processing methods. I am thankful for Dr. Fabbri's advice regarding the implementation of data mining methods and the visualization platform. I would like to thank Dr. Beauchamp for his advice on the real-world information needs of surgical oncologists and their patients.

I would also like to thank the Department of Biomedical Informatics (DBMI) and Medical Scientist Training Program (MSTP) at the Vanderbilt University School of Medicine. The DBMI provided a great environment for my development as a researcher and informaticist. I especially would like to thank the faculty in the department for always being willing to provide career and research advice and serving as a sounding board for various ideas. I would also like to thank the staff and leadership of the department including Rischelle Jenkins and Drs. Cindy Gadd and Kevin Johnson. The MSTP provided a great platform for my clinical, research, and leadership development since my arrival at Vanderbilt. I would like to thank the staff and leadership team at the MSTP who have developed a great culture including Lindsay Meyers and Melissa Krasnove as well as Drs. Terry Dermody and Chris Williams. I am also grateful for the funding that makes my work possible, including the Vanderbilt MSTP T32 and CTSA TL1.

Finally, I would like to extend my gratitude to my family and friends. I appreciate the friendship and camaraderie of my fellow DBMI, MSTP, and medical school classmates and colleagues as we journeyed across the different paths together. I would also like to thank my family for their constant source of support throughout. I especially would like to thank my parents for their lifelong guidance. I would also like to thank my sister (who arrived at Vanderbilt, completed her undergraduate training, and left to get a "real job" all during my time here) for her support.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ACR ..... American College of Radiology

ACS ..... American Cancer Society

AJCC..... American Joint Committee on Cancer

AMA.....American Medical Association

ASCO.....American Society of Clinical Oncology

BCS .....Breast Conserving Surgery

CAP .....College of American Pathology

CoC .....Commission on Cancer of the American College of Surgeons

CDC.....Center for Disease Control

CDS ..... Clinical Decision Support

CPT ..... Current Procedural Terminology

EMR ..... Electronic Medical Record

ER ..... Estrogen Receptor

FORDS.....Facility Oncology Registry Data Standards

HER2 .....Human Epidermal Growth Factor Receptor 2

ICD-9-CM ..... International Classification of Diseases, Ninth Revision, Clinical Modification

IOM ..... Institute of Medicine

IRB .....Institutional Review Board

JCAHO.....Joint Commission on Accreditation of Healthcare Organizations

KBTA .....Knowledge Based Temporal Abstraction

NAACCR .....North American Association of Central Cancer Registries

NAPBC ..... National Accreditation Program for Breast Centers

NCCN ..... National Comprehensive Cancer Network

NCI ..... National Cancer Institute

NPCR ..... National Program of Cancer Registries

NQF ..... National Quality Forum

NQMBC ..... National Quality Measures for Breast Centers

PR ..... Progesterone Receptor

QOPI ..... Quality Oncology Practice Initiative

RD ..... Research Derivative

SEER ..... Surveillance, Epidemiology, and End Results Program

SNOMED CT ..... Systematized Nomenclature of Medicine, Clinical Terms

TIRP ..... Time Interval Related Pattern

TNM ..... Tumor, Node, and Metastasis Classification of Malignant Tumors

UMLS ..... Unified Medical Language System

VICC ..... Vanderbilt-Ingram Cancer Center

VUMC ..... Vanderbilt University Medical Center

Chapter 1: Overview

1.1: Abstract

Pathway exploration can support clinical research efforts and shared decision making between providers, patients, and their families. In this work, we extend the Pathfinder temporal abstraction method and developed a corresponding visualization platform to identify treatment event paths. We used manually curated cancer registry data and administrative data to generate a pathway exploration framework to compare data sources, assess clinical quality measures, re-create the analysis of published studies. We additionally developed a decision aid based on the care path visualizations to support patient education and shared decision making. We demonstrated how our exploration framework is generalizable enough to represent the majority of breast cancer quality measures from certain organizations and that we could support replicating previously published studies. Additionally, we developed, received feedback on, and plan to pilot a patient education decision aid to support shared decision making. Future work is needed to develop data driven methods that fit the abstraction framework and can utilize clinical data that is generated during the course of care. This work can help healthcare providers, organizations, and patients make better healthcare decisions and assess performance.

1.2: Chapter Summary

The focus of Chapter 2 is aimed at describing how understanding paths of care can help healthcare overcome its many challenges, specifically in breast cancer care. Section 2.1 reviews

how continuous pathway exploration is important in helping healthcare tackle its financial challenges, keep up with new advancements, and reduce unnecessary variability. Section 2.2 describes the large footprint of cancer care in the healthcare system and the efforts that have been made to monitor the quality of care. Section 2.3 specifically discusses the role that decision-aids can play in supporting decision making of patients and their families as well as give providers more information to make better recommendations. Finally, section 2.4 describes the challenges of achieving near real time pathway exploration and quality measurement and the opportunities of using electronic medical record data to accomplish it.

Chapter 3 focuses on the methods that have been used to track longitudinal clinical care patterns. Section 3.1 reviews the importance of identifying care patterns for continuous quality monitoring. Section 3.2 compares the impact of various clinical data sources can have on the types of resulting analysis. Section 3.3 reviews prior methods used in identifying care patterns in clinical data including temporal abstraction, data visualization, and event sequence mining. Section 3.3 focuses on reviewing the general components of a temporal abstraction system. Section 3.4 addresses the importance and past work of visualization of abstracted pathways. Finally, section 3.5 discusses the hypothesis and aim of this study to develop a scalable pathway exploration framework.

Chapter 4 describes our Pathfinder methodology for generating quality metrics using general temporal abstraction components. Section 4.1 describes where the clinical data used in this study comes from and how it is stored. Section 4.2 provides an overview of the five subtasks of the Pathfinder method. Section 4.3 focuses on the extraction and representation of the clinical data. Section 4.4 describes the development of our interactive pathway visualization and exploration platform. Section 4.5 describes how this platform can be used to support

clinical research efforts by assessing clinical quality. Section 4.6 describes our development of a patient-facing decision aid to support shared decision making.

Chapter 5 presents the results of the Pathfinder methodology. Section 5.1 focuses on characterizing the data that is extracted for our patient cohort. Section 5.2 describes the distribution of treatment event sequences resulting from the temporal abstraction subtasks. Section 5.3 presents the features implemented for the pathway exploration framework and their results. Section 5.4 describes the ability of the exploration framework to assess clinical quality measures and replicate previously published clinical studies. Finally, section 5.5 describes the development of and feedback received on our data-driven decision aid.

Chapter 6 discusses the conclusions that have resulted from this work and the next steps. Section 6.1 reviews this work's contributions to the informatics domain while Section 6.2 is focused on related limitations and future directions. Sections 6.3 and 6.4 covers the same in the clinical domain.

Chapter 2: Pathway Exploration to Support Quality Measurement and Decision Aids-as-a-Service

Complex healthcare decisions often require decision aids to help patients participate in a shared decision making process (Section 2.1). Breast cancer treatment is multimodal and longitudinal and the growing focus on the disease has led to a special emphasis on quality improvement (Section 2.2). Data-driven decision aids can help patients and providers make informed decisions using a raw data source (Section 2.3). Constructing paths of care has been challenging for healthcare organizations due to the irregular nature of clinical data and the differences in the level of abstraction of concepts between the data and the quality metrics. Care paths can be visualized and explored in a data-driven fashion to help patients and providers make better decisions (Section 2.4).

2.1: Imperative for Pathway Exploration

Healthcare delivery is often multi-disciplinary and involves complex clinical decision-making. As healthcare costs rise, value in healthcare should be centered on the efficiency in generating outcomes for a patient's health status, process of recovery, and sustainability of the achieved health (Porter, 2010). As best practices in healthcare evolve, it is also important to track the impact of new clinical knowledge. There have been significant advances in the development of new diagnostics and therapeutics (Collins, 2011), advancement of genomic medicine (Hamburg & Collins, 2010), and use of patient-generated data to personalize care (M. Swan, 2012). Patients are often unable to participate in a shared decision making process with

their provider as they might have difficulty understanding their care plan and might feel intimidated by the decision making process (Knops et al., 2013; Waljee, Rogers, & Alderman, 2007). Decision aids can help providers, patients, and their families make better decisions and we believe we can support an effective shared decision making effort by developing population-based data visualizations.

"Unnecessary variability in clinical care also plays a major role in the quality and value of clinical care. Variability that stems from poor care processes, improper usage of health IT systems, or divergence from clear guidelines should be curtailed; institutional efforts to understand and rectify these issues have been shown to improve the quality and efficiency of care (Ancker et al., 2014; James & Savitz, 2011). However, it is still important to allow for flexibility from standardized clinical guidelines to enable patient choice, socioeconomic considerations, and clinical complexity. In assessing clinical quality, utilization of shared decision making between patients and providers should be viewed as a positive, even if the patient makes a choice that deviates from the standard path (Oshima Lee & Emanuel, 2013; Quill & Holloway, 2012). Additionally, in cases where patients have multiple comorbid conditions, providers must be empowered to focus on the patient as a whole rather than apply multiple clinical guidelines (Boyd et al., 2005; Tinetti, Fried, & Boyd, 2012). While many have concentrated on the role of geography in the variability of patient care, efforts to reform clinical decision making should still be targeted at where decisions are made: healthcare providers, organizations, and networks (Newhouse et al., 2013). Given both the necessary and unnecessary variability in healthcare, it is important to understand clinical care paths from the perspective of the individual patient case" (Atreya, 2015).

Breast cancer care is a microcosm of the many challenges facing healthcare due to its multimodal and longitudinal treatment as well as its growing survivor population. There are three primary modes of treatment that are administered based on the stage, tumor biology, and patient's preferences. In the curative setting (stages I-III), surgery is used to excise the tumor and regional lymph nodes. Surgical options include breast conserving surgery or mastectomy with or without reconstruction. Radiation therapy is used as an adjuvant therapy to surgery to provide local control at the tumor site through use of ionizing radiation to damage of tumor. Systemic drug therapy treats both the breast and the rest of the body. There are three types of systemic therapy used to treat breast cancer: chemotherapy, hormone therapy for hormone receptor positive disease, and anti-HER2 therapy for HER2 positive disease. The choice and sequencing of these complex multi-modal treatments depends upon both tumor and patient features.

In oncology care, clinical pathways have traditionally been used to define evidence-based care plans based on specific patient and tumor characteristics. Pathways began to be used after the 2003 Medicare Prescription Drug, Modernization, and Improvement Act as a method to balance cost and quality of care. Some pathway systems exist as standalone web portals where providers enter patient information and the system generates clinical recommendations. Pathways have also been integrated into electronic medical record (EMR) systems in order to reduce errors, increase efficiency, and help streamline reimbursement procedures (Gesme & Wiseman, 2011; Zon et al., 2016). Studies have demonstrated that pathways have reduced costs while at least maintaining the quality of care (Hoverman et al., 2011; Neubauer et al., 2010). On the other hand, pathways are sometimes criticized due to the transparency in how they are developed, the administrative burden of a growing number of

pathways, and that they might be too 'cookie-cutter' (Gesme & Wiseman, 2011; Kinsman, Rotter, James, Snow, & Willis, 2010; Zon et al., 2016).

While clinical pathways use protocols to define the care path, data-driven exploration of sequences of diagnostic and treatment events can allow patients, providers, and clinical researchers to explore care paths on their own. Understanding care paths can be aided through effective data visualization and allowing technology to augment human reasoning (Friedman, 2009; Miller & Masarie, 1990; Spence, 2006).

2.2: Importance of Quality Metrics in Cancer Care

"Breast cancer is the most common cancer among women in the United States with an estimated 232,670 new cases diagnosed and 40,000 breast cancer related deaths in 2014. Women in the United States have a 12.3% lifetime risk of developing breast cancer, but benefit from a very high median 5-year survival rate of 89.2% (Siegel, Ma, Zou, & Jemal, 2014). Breast cancer accounts for 14% of all cancer cases in the US, and 13.2% of total cancer costs. With the increasing number of survivors, including 3.46 million women in 2010, and an aging population, the estimated 2010 cost of $16.5 billion is expected to increase 24% over the following decade (Mariotto, Yabroff, Shao, Feuer, & Brown, 2011). Breast cancer will continue to be one of the primary areas of care delivery, especially as the number of survivors continues to grow" (Atreya, 2015).

"With advances in breast cancer research and the trend towards shared decision-making, there is growing tension between personalized and pathway driven care. Evidence behind the NCCN guidelines has shown that only 6% of the guidelines are based off category I

evidence, such as randomized controlled trials (RCTs). In breast cancer, staging guidelines have 100% of the content based on level IIA evidence (lower level evidence with consensus); initial therapy guidelines have 42% of content based on level I, 42% on level IIA, 11% on level IIB (lower level evidence without uniform consensus but no major disagreement), and 5% on level III evidence (major disagreements); salvage therapy guidelines have 100% of the content based on level IIA evidence; surveillance guidelines have 67% of the content based on level IIA and 33% on level IIB evidence (Poonacha & Go, 2011). The appropriate level of adherence to guidelines in comparison to deviation from the pathway in order to personalize care is uncertain given the varying levels of evidence behind the guideline (R. C. Chen, 2013). As quality reporting becomes tied to reimbursement and the value of personalized diagnostics and therapeutics are determined by payers (Weldon, Trosman, Gradishar, Benson, & Schink, 2012), the strength of evidence behind guidelines is another important factor to determine the balance between pathway-based and personalized care" (Atreya, 2015).

"Since the 1970's, the number of breast cancer centers delivering complex care has increased. In order to externally validate the quality of these centers, accreditation bodies have been instituted that enable breast centers to voluntarily participate in trusted quality assurance programs. External accreditation began with the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) in 1951 and a number of cancer-specific accreditation bodies have been instituted since three major IOM reports were released in 1999-2001 (Edge, 2013). Breast cancer specific accreditation programs include the National Accreditation Program for Breast Centers (NAPBC) (Burgin, 2010; "National Accreditation Program for Breast Centers (NAPBC)," n.d.) and the National Quality Measures for Breast Centers (NQMBC) (Kaufman et al., 2010; "National Quality Measures for Breast Centers (NQMBC)," n.d.). Many professional

organizations also have breast cancer quality components as part of their accreditation including The Quality Oncology Practice Initiative (QOPI) (Goldwein & Rose, 2007; Peterson, 2012; "Quality Oncology Practice Initiative (QOPI)," n.d.), National Quality Forum (NQF) Breast Quality Measures ("National Quality Forum (NQF)," n.d.; Thompson et al., 2012), the College of American Pathology (CAP) ("College of American Pathology (CAP)," n.d.), the American College of Radiology (ACR) ("College of American Radiology," n.d.), and the Commission on Cancer (CoC) ("Commission on Cancer," n.d.; Minami et al., 2016). The most common method of accreditation among these organizations is through quality reporting. Complying with the quality audits is costly and time consuming, with overall hospital accreditation and licensure costs estimated to be $8.6 billion annually (Conover, 2012). Because of the high costs, the quality reporting is periodic and not continuous which limits its value. In addition, there are mixed results in analyzing the success that accreditation programs have on improving the quality of care (Greenfield & Braithwaite, 2008; Merkow, Chung, Paruch, Bentrem, & Bilimoria, 2014). The growth of accreditation for cancer care has increased the pressure on cancer centers to track quality and on the accreditation programs to demonstrate their value in improving care" (Atreya, 2015).

Quality metrics play a major role in assessing quality for accrediting organizations and for new quality and population-oriented payment models. Three major categories of quality metrics include clinical metrics that assess the quality of clinical decision-making, process metrics that measures the operational execution of a clinical decision, and outcome metrics that evaluates the clinical outcome for the patient. In this study, we chose to focus on clinical metrics that involved patterns of care. These metrics assess the patterns of care with regard to systemic therapy, surgery, and radiation therapy usage. Quality metrics are a vital component

for improving the quality of breast cancer patients.

A significant effort to record cases of cancer began with the National Cancer Act of 1971 and the establishment of the Surveillance, Epidemiology, and End Results Program (SEER) at the National Cancer Institute (NCI) (Warren, Klabunde, Schrag, Bach, & Riley, 2001). SEER collects data from select regions around the country. In order to expand this collection effort, the Cancer Registries Amendment Act of 1992 established the National Program of Cancer Registries (NPCR) at the Center for Disease Control (CDC) to expand state cancer registry programs. SEER and NPCR are considered population-based registries. On the other hand, the National Caner Data Base (NCDB), which is managed by the Commission on Cancer (COC) via the American College of Surgeons (AcoS) and the American Cancer Society (ACS), is a hospital-based registry. Participation in the NCDB is required in order to be designated as a COC-approved cancer center(J. Swan et al., 1998; P. Wingo & Howe, 2005).

The learning health system extends past the system of cancer registries and aims to continuously learn best practices from previously seen patients in order to inform the care of patients in the future. To accomplish this, clinical & patient generated data should be continuously collected, analyzed, and utilized to generate new knowledge. Newly developed knowledge should then be implemented in clinical care that will generate additional data that can be evaluated in order to generate additional hypothesis (Abernethy et al., 2010; Etheredge, 2007, 2014). One example of an effort to put the learning health care system concept into practice is ASCO's CancerLinQ program (Sledge, Miller, & Hauser, 2013). It is modeled as the successor to the Quality Oncology Practice Initiative (QOPI) – it aims to collect clinical data in real time and analyze & compare it to guidelines. CancerLinQ will, for example, collect comprehensive EMR data (i.e., demographics, billing codes, notes, history, laboratory data,

medications). This data will be used to deliver real-time clinical decision support to providers and to update clinical practice guidelines. CancerLinQ is ASCO's implementation of a learning health system as it aims to collect, aggregate, and analyze clinical data to generate new knowledge while feeding insights back to providers(Sledge, Hudis, et al., 2013; Sledge, Miller, et al., 2013).

2.3: Decision Aids in Breast Cancer Care

Breast cancer patients face difficult decisions about their care without a full understanding of their options. Women often have a choice between a breast conserving surgery (BCS), where only the tumor is removed, and mastectomy, where the entire breast is removed. It is especially important to understand surgical options, as the procedures are often irreversible. While some patients prefer to have their doctor take the lead decision making role, many patients want to participate in a shared decision making process. Decision aid tools can help provide the patient with easy to understand information on potential options (Knops et al., 2013; Waljee et al., 2007).

Traditional decision aid studies in breast cancer care have utilized brochures, booklets, audiotapes, and videotapes to assess items such as the patients' knowledge, final treatment decision, patient satisfaction, and decisional conflict/regret. Decision aids focused on early stage breast cancer tended to help fortify patient knowledge, reduce decisional regret, increase patient satisfaction, and improve doctor-patient communication. Additionally, use of a decision aid for surgical decisions tended to lead patients to choose the more conservative and less invasive option (Knops et al., 2013; Waljee et al., 2007).

The content in traditional decision aids is often written or recorded based on information in previously published studies. Data-driven decision aid tools, instead, directly use actual data to help patients and providers make decisions. There are a variety of these data-driven tools online that pull from a variety of data sources and have differing target audiences. In breast cancer care, "Adjuvant! Online" supports decision making for adjuvant systemic therapy for women with early stage breast cancer (Ravdin, 2001). The survival rates per treatment type are based on data from the SEER cancer registry and various meta-analysis from adjuvant breast cancer clinical trials. This tool is intended to only be used by healthcare providers and requires a login. Providers can review and share a printout of the online tool with the patient. Other tools are broadly available to the public. For example, the NorthShore "What's Going Around" tool shows the fraction of patients with strep throat, influenza, pertussis, pediatric asthma, and gastroenteritis that have been seen at the NorthShore University HealthSystem on a geographic heat map of the Evanston region (Campbell et al., 2015). The data is derived from the NorthShore electronic medical record (EMR) system that is analyzed nightly to identify newly diagnosed patients with the illnesses using a set of logistic regression models. Finally, ProPublica's "Surgeon Scorecard" helps prospective patients review specific surgeon and hospital adjusted complication rates based on 30-day readmissions across eight specific surgical procedures (Pierce & Allen, 2015). This tool is based on Medicare billing data from 2009-2013. This tool only showed the adjusted complication rate and surgical volume per surgeon, but does not show any additional information on the patient population or adjustment process per-surgeon.

2.4: Challenges and Opportunities for Pathway Exploration in Cancer Care

Generating explorable care paths from medical record data is a challenge due to the unstructured and irregular nature of the content. "Curated databases, such as cancer registries, contain manually structured, standardized data that makes it easier to build care paths. However, the maintenance of the data requires significant time and manpower investments. In addition, cancer registries are often missing important information as they do not include the entire clinical population due to their inclusion criteria and do not contain all the clinical data due to the limited reporting requirements. Even with structured clinical data, understanding the information at the proper level of abstraction is a major challenge. However, data stored in the registry or medical record is often not at the level of abstraction posed by the clinical question. NAACCR, SEER, the NCI Thesaurus, and others work to meet this challenge by developing hierarchical ontologies and dictionaries to represent clinical data (Bilimoria, Stewart, Winchester, & Ko, 2008; de Coronado, Haber, Sioutos, Tuttle, & Wright, 2004; "NCI Thesaurus," n.d.). Understanding the abstraction of clinical data is important in representing terms at the proper level on a hierarchy as well as representing terms that represent a pattern of care over time" (Atreya, 2015).

"The need for variable and dynamic abstraction is demonstrated by the breast conserving surgery event which is a part of both the re-excision rate and radiation therapy quality metrics. First, the "is-a" relationship allows a lumpectomy event to also be represented as a breast conserving surgery and surgery event. Maintaining this "is-a" hierarchy allows for clinical concepts to be vertically abstracted to a certain level on the hierarchy depending on the quality metric being evaluated. Figure 1 represents this vertical abstraction from a subsection of

the NCI Thesaurus. Second, the sequence of surgical events over time also determines how the event and a patient's course of care should be classified. For example, if two breast conserving surgeries occur during the course of care, the second surgery could be understood as a re-excision event because a second invasive surgery was required. This requires that a sequence of events be defined through a set of horizontal abstractions that can represent the temporal nature of event sequences. Figure 1 demonstrates the set of horizontal abstractions that help define the consolidation of surgical events into representing either a breast conserving surgery or mastectomy based on the presence of a mastectomy in the surgical sequence. These vertical and horizontal abstractions are necessary to organize patient cohorts by treatment event sequences in order to identify patient pathways" (Atreya, 2015).

Finally, delivering pathway exploration to stakeholders effectively in a dynamic fashion is a major challenge. The development of these care paths is only useful if they have a tangible impact on the quality of care. "This will require study of how best to visualize and communicate this information to providers, patients, and administrators. As cancer providers continue to adopt electronic medical record (EMR) systems, it may become possible to track clinical care in real time as information is aggregated digitally" (Atreya, 2015). New algorithms will be required that can learn from data, identify patterns, and abstract the data accordingly. In addition, we will need new ways to communicate this information to healthcare stakeholders to allow providers, patients, and their families to learn from every clinical encounter.

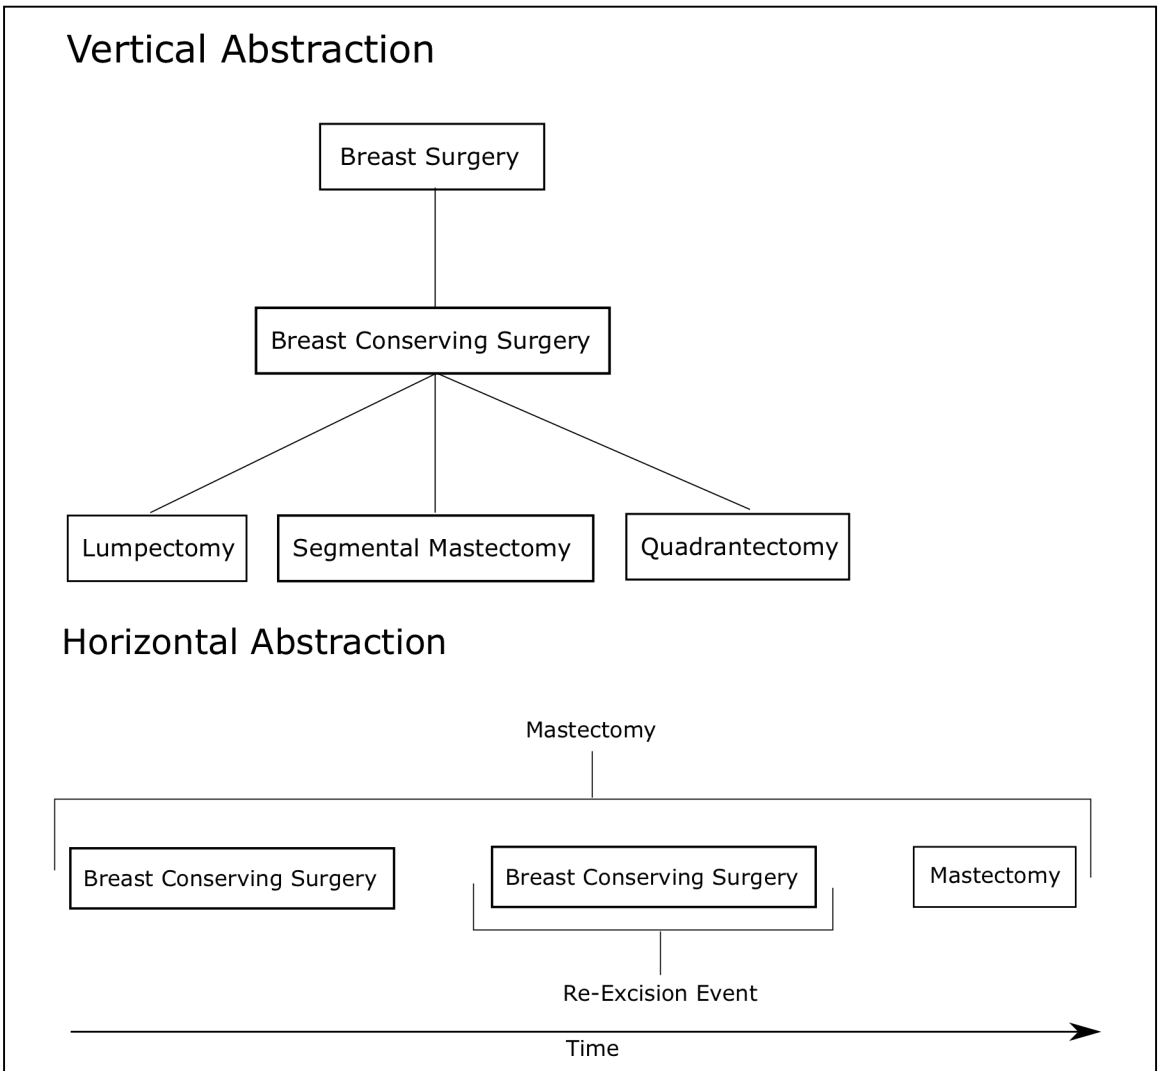

Figure 1: Example of vertical and horizontal abstraction for linking clinical data with quality metrics

Top: Vertical abstraction for the concept Breast Conserving Surgery demonstrates how concepts can be organized hierarchically using "is-a" relationships.

Bottom: Horizontal abstraction that demonstrates how the interpretation of events can be determined by the temporal sequence of events. In this case, the second breast conserving surgery can be considered a re-excision event while the full set of events is defined by the resulting mastectomy.

Chapter 3: Tracking Clinical Care Patterns

Pathway exploration requires multiple subtasks to convert raw clinical data to care paths and patterns (Section 3.1). Care pathways have the potential to be developed from cancer registry and electronic medical record data (Section 3.2). Temporal abstraction has the potential to enable this process as it allows for the representation of clinical data at varying levels of granularity based on its usage. The use of formal, knowledge-based models to represent temporal concepts with clinical data began in the 1980s and have since grown (Section 3.3). Visualization methods have allowed clinical experts to identify clinical patterns from the abstracted information (Section 3.4). This study will aim to develop a scalable framework to identify and visualize patterns of care at multiple levels of abstraction and use it to support a pathway exploration framework (Section 3.5)

3.1: Imperative for Tracking Clinical Care Patterns Over Time

Pathway exploration requires the ability to dynamically abstract clinical data and identify complex clinical care patterns. Clinical data entered into the medical record system is often not at the level of vertical abstraction necessary to answer various clinical questions. Furthermore, clinical data can be missing, unevenly spaced over time, and unstructured. As a result, methods to identify the optimal set of abstractions for clinical data elements are vital for allowing pathways to be explored in the most clinically relevant fashion (Figure 2). "Additionally, many quality metrics require an understanding of the patient's treatment pathway. Effective horizontal abstraction and pattern recognition of clinical events is important

in understanding the patient's treatment course. Methods that can incorporate clinical context across time to aid in clinical data interpretation are required to abstract clinical data horizontally and vertically to generate care paths" (Atreya, 2015).

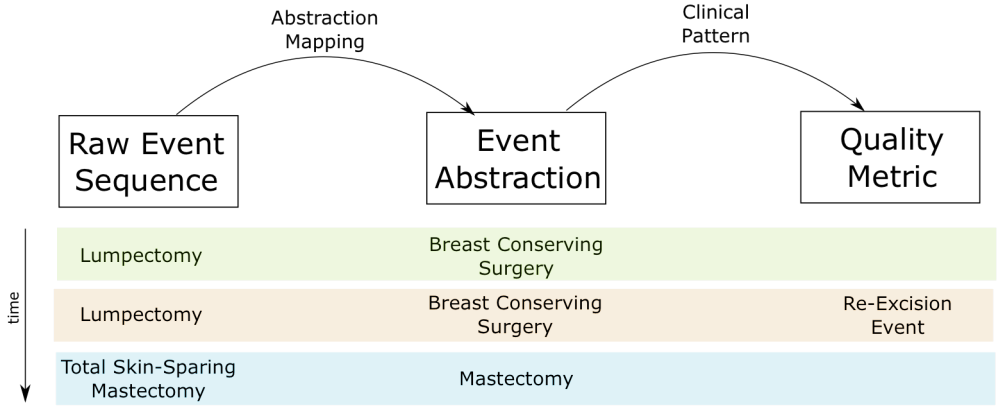

Figure 2: Converting raw event sequences to information for pathway exploration

Pathway exploration requires that raw event sequences be abstracted to understand the series at the proper level of granularity necessary to generate the desired quality metric.

3.2: Impact of Clinical Data Sources of Extracted Care Paths

Two primary sources of clinical data for cancer patients include the electronic medical record (EMR) and the cancer registry. Medical record data is generated through the normal course of care by healthcare providers. While it is often more detailed, it is also often messy, irregular, and can be prone to missing data on care that took place at external facilities. Cancer registries, on the other hand, are manually curated, structured using a customized encoding scheme, and often fill in care delivered externally. Information contained in registries is less detailed than what is stored in the medical record. Finally, while the population of patients in a

cancer registry at a single institution will be smaller than the full cancer patient cohort, cancer registries can be aggregated nationwide. As a result, cancer registries can cover a majority of the United States population (Figure 3).

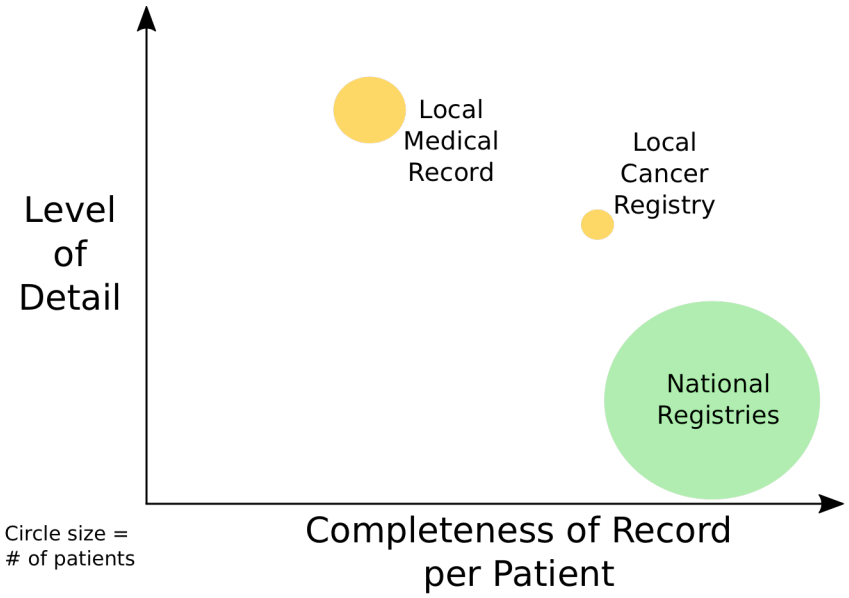

Figure 3: Tradeoffs in detail, completeness, and size between data sources

Local medical record data, while rich in detail with laboratory results, clinical notes, and other granular details, only contains information on the subset of clinical care that their patients receive at that institution. The local cancer registry curates a more complete picture of a cancer patients' clinical care but only across a select number of data types. The national registry can compile the largest dataset based on smaller registries.

The National Coordinating Council for Cancer Surveillance (NCCCS) is a consortium of the ACS, AcoS, CDC, NCI, NCRA, and the North American Association of Central Cancer Registries (NAACCR) that was formed in 1995 to facilitate the process of coordinating data standards and collection requirements. NAACCR helps develop standards for population-based registries while the COC helps develop standards for hospital-based registries. As a result of coordinating

efforts, there has been a general standardization in the data collection process (J. Swan et al., 1998; P. Wingo & Howe, 2005). Any neoplasm with a behavior code of 2 or 3 (*in situ* or malignant) is required to be reported for the NCDB, NPCR, and SEER using ICD-O-3 (ICD-Oncology, 3rd ed.) with the exception of squamous cell and basal cell carcinoma of the skin and carcinoma *in situ* of the cervix uteri. Additionally, the NCDB, NPCR, and SEER require the reporting of all non-malignant primary intracranial and central nervous system tumors. The SEER and NPCR population-based registries will consolidate data from the various reporting sites. The NCDB, as a hospital-based registry, utilizes a class of case designation to define analytic and non-analytic cases. Analytic cases indicate a hospital's primary role in managing the patients' cancer and must be fully abstracted (including treatment and outcome). Non-analytic cases only require an indication of why a patient was seen at the institution and any additional compliance with locally set standards ("NAACCR Chapter III: Standards for Tumor Inclusion and Reportability," n.d., "NAACCR Chapter X: Data Dictionary," n.d.; J. Swan et al., 1998; P. Wingo & Howe, 2005).

The cancer registry is updated on a regular basis as cancer cases are identified (casefinding), reported, and abstracted by cancer registrars. Certified Tumor Registrars (CTRs) are cancer registrars who have achieved certification via a training and continual education program which was developed by the National Cancer Registrars Association (NCRA) (J. Swan et al., 1998). Data is collected on patient demographics, cancer stage, tumor pathology, co-morbidities, treatments, and vital status so that survival outcomes can be measured (i.e., 5-year survival). The registry staff use the NAACCR cancer registry data dictionary, the Facility Oncology Registry Data Standards (FORDS) manual from the CoC, and the American Joint Committee on Cancer (AJCC) Tumor-Node-Metastasis (TNM) staging standards ("NAACCR Chapter X: Data Dictionary," n.d.; J. Swan et al., 1998).

State cancer registries are managed by the NPCR via the CDC and are a population-based registry. The CDC supports cancer registries in 45 states, Washington DC, and three US territories, covering 96% of the US population. Given the large population captured in the NPCR (as well as each individual state registry) and the abstracted and structured nature of the data acquired, the NPCR data can be utilized to assess national and state-wide trends over time with regard to cancer diagnosis, treatment, and outcome. For example, the NPCR is used (along with SEER) to generate the annual United States Cancer Statistics (USCS) that compiles incidence data. The primary benefit of the NPCR is as a population-based registry that has almost complete national coverage. A weakness of the NPCR is the lack of information collected on recurrence. In fact, only the NCDB collects recurrence data (but is not publically available or reported on due to concerns on completeness). Another possible weakness (depending on the goal of the study) is the relative underrepresentation of minority populations compared to SEER (In et al., 2014; P. A. Wingo et al., 2003).

In comparison to cancer registries, there are many data resources typically stored in an electronic medical record (EMR) system that can be used for phenotyping.(Denny, 2012) Administrative data includes billing data such as ICD-9-CM/ICD-10-CM codes (International Classification of Diseases) and CPT codes (Current Procedural Terminology). ICD-9-CM covers codes used to bill for the management of diseases, identification of symptoms, as well as a few procedures. ICD-9-CM is the US version of the ICD-9 terminology developed by the World Health Organization. CPT is used to code for procedures and other clinical services. CPT was developed by and is maintained by the American Medical Association. When used for phenotyping (for a disease state), CPT codes will have a high specificity and low sensitivity. ICD-9-CM codes, on the other hand, will have low specificity but high sensitivity.

Clinical medical record data includes laboratory data, orders/reports from ancillary

services, medications, and clinical notes. There will often be multiple laboratory results and vital signs recorded in the EHR per patient in a longitudinal fashion. Vital signs will often be recorded as a structured value while laboratory results can be reported in a structured or unstructured form (will be discussed later). Pathology, imaging, and other testing orders and reports convey specific information on the type and results of a clinical test. These reports will often contain both structured and unstructured content in free text (i.e., data might have been partially inputted using a structured reporting tool). Medication data indicates that a clinician decided to prescribe a medication. Additionally, changes in medication list (i.e., medications, dosages) can be tracked over time. Medication data can be found in a variety of sources depending on the EHR being used (i.e., structured computerized provider order entry, bar code administration, structured medication reconciliation). Clinical notes that are typed into EHRs can be used to understand how healthcare providers described the patients' history of present illness, past medical/surgical history, social and family history, physical exam, and assessment/plan at multiple points in time. These notes are often unstructured (unless generated through a structured coding user interface).

3.3: Organizing Clinical Data to Represent Treatment Patterns

"Prior methods for tracking clinical care patterns have included heuristic, knowledge based, and probabilistic approaches. Temporal abstraction methods aim to use knowledge-based and heuristic processes to organize clinical events at a higher level and elicit useful information from the abstractions. Data visualization has been utilized for empowering clinical experts to identify patterns from longitudinal clinical data of both individual patients and clinical populations. Finally, data mining approaches, including event sequence mining methods, can

aid in the discovery of temporal associations from clinical data. Temporal abstraction methods are valuable in representing data on multiple levels; they can help improve human cognition through the incorporation of data visualization methods and can be made more scalable through the use of data mining methods" (Atreya, 2015).

Methods for temporal representation and reasoning in medicine have been extensively reviewed (Adlassnig, Combi, Das, Keravnou, & Pozzi, 2006; Augusto, 2005; Combi & Shahar, 1997; M. G. Kahn, Fagan, & Sheiner, 1991; Orphanou, Stassopoulou, & Keravnou, 2014; Stacey & McGregor, 2007). Of those, the Knowledge Based Temporal Abstraction (KBTA) method developed by Yuval Shahar and colleagues, remains the central foundation for a knowledge-level framework for formalizing the requirements for an abstraction ontology and methodology (Y Shahar, Tu, & Musen, 1992; Yuval Shahar, 2013; Yuval Shahar & Musen, 1996). The representation of and relationships between temporal events and intervals were originally formally defined in the 1980s to represent time series data in the field of artificial intelligence (Allen, 1984; Kowalski & Sergot, 1989; Mcdermott, 1982; Shoham, 1987). Original clinical knowledge based systems such as MYCIN (Shortliffe, 1977) and Internist-I (RA, HE, & JM, 1982) focused on representing clinical domain knowledge and used their symbolic structures to represent temporal features along with the clinical ontology in a post-coordinated manner without dynamically abstracting the temporal data. The first system to utilize temporal abstraction was Fagan's ventilator management system that tracked context-specific rules for clinical parameters (Fagan & Kunz, 1984). Early systems that handled time series data did not use a temporal ontology and used simple abstraction hierarchies (i.e., summarization program (Downs, Walker, & Blum, 1986), IDEFIX (de Zegher-Geets, Freeman, Walker, Blum, & Wiederhold, 1988)) or maintained simple relationships (i.e., TCS (Russ, 1995)). The TOPAZ

system had a temporal interpretation scheme with a management and query system, however it was domain-specific and lacked generalizability (M. G. Kahn, 1991; M. G. Kahn, Fagan, & Sheiner, 1991; M. G. Kahn, Fagan, & Tu, 1991; M. Kahn, Tu, & Fagan, 1991). The TrenDx system used specific pre-defined clinical event patterns and had the ability to fit partial patterns, while avoiding a more data-driven approach (I. J. Haimowitz & Kohane, 1993; I. Haimowitz & Kohane, 1993; Kohane & Haimowitz, 1993). Temporal abstraction systems have advanced from developing temporal representation structures to developing the ability to organize medical record data temporally (Atreya, 2015).

"The original KBTA implementation was the RESUME system that implemented the temporal abstraction platform and has been tested on a variety of clinical domains including oncology, AIDS, and insulin-dependent diabetes (Y Shahar & Musen, 1992; Yuval Shahar & Musen, 1996). An ecosystem of tools was developed around the RESUME system in order to operationalize it: CAPSUL is a temporal pattern language that was developed and integrated with RESUME to allow complex pattern creation and querying (S Chakravarty & Shahar, 2001; Shubha Chakravarty & Shahar, 2001); EON extended its capabilities to therapy planning (Musen, Tu, Das, & Shahar, 1996); TZOLKIN contained the database architecture and enabled the querying of the data based on their abstraction goals (Nguyen, Shahar, & Tu, 1999); Asbru was developed as a method of representing skeletal guideline plans (Miksch, Shahar, & Johnson, 1997) which was a part of the Asgaard project involved with utilizing those plans in the clinical domain (Yuval Shahar, Miksch, & Johnson, 1998); ALMA (Balaban, Boaz, & Shahar, 2003) and IDAN (Boaz & Shahar, 2005), the updated versions of RESUME and TZOLKIN, use a distributed architecture and can be linked to various knowledge bases and ontologies; Momentum allows the system to handle streaming data as information is entered into the clinical record (Spokoiny

& Shahar, 2003). A large complement of tools was shown to be necessary to implement, maintain, and scale a temporal abstraction system in the clinical setting" (Atreya, 2015).

"Data mining and statistical approaches can also help identify treatment patterns from longitudinal clinical data. Association rule mining is a method for identifying association patterns between categorical variables and was often used in analysis of customer transaction data (Rakesh Agrawal, Imieliński, & Swami, 1993). Event sequence mining, an extension of association rule mining, has aimed since its introduction (R. Agrawal & Srikant, 1995) to identify temporal event patterns based on time stamped data (Bellazzi, Ferrazzi, & Sacchi, 2011; Bellazzi, Sacchi, & Concaro, 2009). In addition to customer transaction data, event sequence mining methods have been applied to identifying plan failures and network alarms and research has focused on developing more efficient algorithms (Zaki, 2001). These methods have also been applied to clinical data in KarmaLego (Moskovitch & Shahar, 2009) (part of RESUME KBTA ecosystem) and ChronoMiner (Raj, O'Connor, & Das, 2007) (ontology-based pattern mining), as well as to administrative data (Norén, Hopstadius, Bate, Star, & Edwards, 2009). KarmaLego uses event sequence mining methods in the context of its KBTA framework to identify time interval related patterns (TIRPs). It was demonstrated on a set of diabetes patient data and was used to discover and classify patient subgroups. Chronominer is another pattern mining system that searches for patterns simultaneously at multiple levels of abstraction and was tested on HIV patient data to assess genetic mutations acquired during therapy. Event sequence mining can be an effective method for identifying common event patterns in clinical data in the context of a temporal abstraction framework" (Atreya, 2015).

Generating explorable care paths is challenging and often requires the abstraction of clinical data and the use of event patterns. Temporal abstraction and data mining methods will

be necessary to align raw clinical data with longitudinal pathways. Breast cancer treatment is multimodal and variable. As a result, efforts to study clinical quality have been slow and required manual review. We plan to develop a temporal abstraction framework to develop a pathway exploration framework for breast cancer care.

3.4: Visualizing Paths of Care

"Learning patterns of care from abstracted clinical data can be improved through effective data visualization and allowing technology to augment human reasoning (Friedman, 2009; Miller & Masarie, 1990; Spence, 2006). A 1994 article was one of the earliest to propose using graphs to summarize patient data from multiple sources (Powsner & Tufte, 1994). Since then, the two most published tools for temporal data visualization have been the LifeLines (Plaisant, Milash, Rose, Widoff, & Shneiderman, 1996)/LifeFlow (Krist Wongsuphasawat, 2011) systems by Plaisant and Shneiderman as well as the KNAVE (Yuval Shahar, Goren-Bar, Boaz, & Tahan, 2006)/VISITORS (Klimov, Shahar, & Taieb-Maimon, 2010) tools by Shahar which are a part of the RESUME KBTA ecosystem. Plaisant and Shahar both initially focused on generating graphs for the clinical data of a single patient and have progressed to visualizing the trends of care across thousands of patients. These and other systems have been used to help clinical experts more effectively understand clinical records of individual patients and patterns of care across clinical populations (K. Wongsuphasawat & Gotz, 2012; Zhang, Wang, Ahmed, & Ramakrishnan, 2013) " (Atreya, 2015).

Patient level visualizations can be helpful in summarizing a complex medical record. For example, the LifeLines system represents events of variable-length duration on a timeline so

that a user can see how the occurrence of various events are related. While patient level timelines allow the user to identify patient level patterns, population level visualization systems such as LifeFlow allow users to filter on multiple patient records in a compact fashion. Scalable population-based visualization tools need to be able to represent variable sized populations from thousands to millions of patients using a similar representation (Plaisant et al., 1996; Krist Wongsuphasawat, 2011). One challenge for developing scalable visualizations is the number of potential values for specific features. For example, there are many kinds of unique surgical procedures; these can be organized into a hierarchical structure. By using the higher levels of representation, visualizations can show more information more compactly (Guerra-Gomez, Pack, Plaisant, & Shneiderman, 2013).

Of fifteen articles between 1996 and 2013 that discuss temporal clinical data visualization, 13 focused on clinical decision support while only two focused on quality improvement (West, Borland, & Hammond, 2014). With a growing number of quality metrics and increasing cost of care, data visualization across a medical record system can allow for more effective analysis of the quality of clinical practice ("Advancing Meaningful Use: Simplifying Complex Clinical Metrics Through Visual Representation," 2010). Data visualization can help clinical experts to identify treatment patterns across clinical populations and effectively deliver quality metrics to various stakeholders.

In this work, we aggregate and display the abstracted treatment paths across our breast cancer patient population(Krist Wongsuphasawat, 2011). Aggregation is necessary to enable scalable visualization of large patient cohorts and avoid trying to display a large number of individual patient records in a confined visual space. This aggregation step involves the grouping of events across treatment paths to find common prefix, root, and/or suffix sequences (Figure

4). The visualization step is focused on using color-coded event bars to represent events where the color of the bar represents the event type, the height of the bar represents the fraction of patients with that event in the sequence, and the position of the bar represents the position of the event in the sequences treatment events. Using the fraction of patients is scalable as a constant visual space can represent 100% of the population. The LifeFlow visualizations (Krist Wongsuphasawat, 2011) use prefix aggregates and are based on the Icicle Tree diagram. This has the potential to lead to a complex tree of sequences when the treatment paths are long and highly variable(Monroe, Lan, Lee, Plaisant, & Shneiderman, 2013). The Sankey diagram aims to further group sequences by representing prefix, root, and suffix aggregate patterns using a similar design principle (Riehmann, Hanfler, & Froehlich, 2005). This leads to greater consolidation across patients along the full length of the treatment path as event sequence aggregates can both diverge and converge.

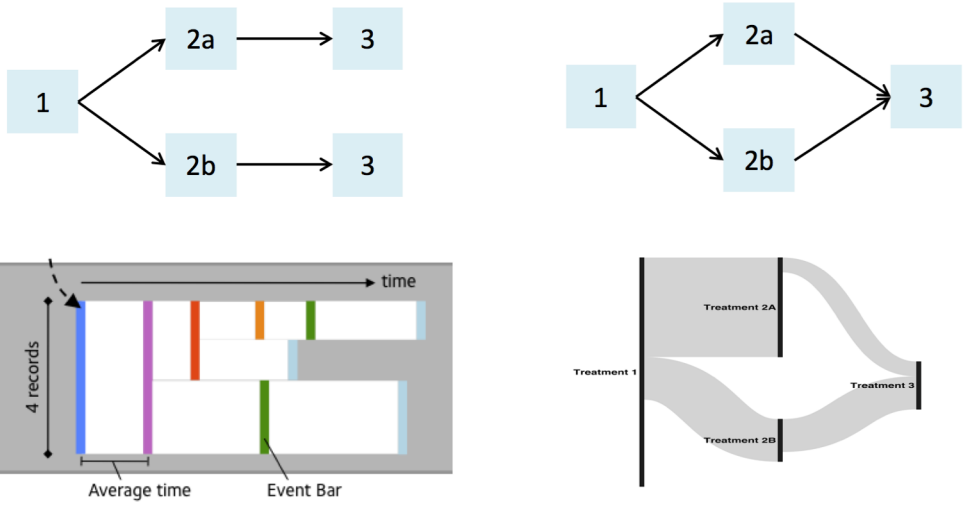

Figure 4: Visualizing aggregate paths with icicle tree (L) diagrams and Sankey (R) diagrams

Visualizing aggregate paths involves representing prefix, root, and suffix aggregate patterns. Both the icicle tree (left) and Sankey (right) diagrams handle prefix aggregates (1 → 2a) by displaying a diverging path. The icicle tree and Sankey diagrams handle suffix aggregates (2a/2b → 3) differently as icicle tree diagrams do not converge while Sankey diagrams do. As a result, Sankey diagrams can be more compact.

3.5: Hypothesis and Aims

Our objective is to develop a scalable framework to identify and visualize patterns of care at multiple levels of abstraction and use it to support a pathway exploration framework. We hypothesize that a framework that consists of vertical and horizontal abstraction methods can help generate data-driven care paths and develop interactive path visualizations to inform patients and providers.

Aim 1: Develop a scalable framework to extract and abstract breast cancer event data to develop a pathway exploration framework

Aim 2: Evaluate the quality of care using a pathway exploration framework

Aim 3: Visualize care paths to deliver decision aids-as-a-service

Chapter 4: The Pathfinder Framework

This chapter describes the Cancer Pathfinder framework and its application to supporting pathway exploration. The data for this study is derived from the National Cancer Database and the Vanderbilt University Medical Center (VUMC) cancer registry (Section 4.1). The Cancer Pathfinder method consists of five sub-tasks including data extraction, data standardization, vertical abstraction, horizontal abstraction, and pathway exploration (Section 4.2, 4.3). We developed a tool to dynamically and interactively explore care paths of breast cancer patients (Section 4.4). We next demonstrated how the pathway exploration tool could support quality measurement (Section 4.5). Finally, we developed a data-driven decision aid that uses interactive care path visualizations to help patients engage in a shared decision making process (Section 4.6).

4.1: Clinical Setting and Patient Data Sources

This study primarily used data collected from the cancer registry at the National Cancer Database (NCDB) and the Vanderbilt-Ingram Cancer Center (VICC). Additionally, data was collected from clinical information systems at Vanderbilt University Medical Center (VUMC) and Vanderbilt-Ingram Cancer Center (VICC). The identifiable patient data from Vanderbilt used in this study included information on demographics, treatments, billing codes, outcomes, and providers. The de-identified patient data from the NCDB included demographic and treatment information. This study has been reviewed and approved by the VUMC Institutional Review Board (IRB) as expedited and minimal risk health sciences study #130957.

The NCDB is an aggregation of cancer registry data from over 1500 Commission on Cancer (CoC)-accredited sites nationwide and is jointly sponsored by the American College of Surgeons (ACoS) and the American Cancer Society (ACS). New patient records in the NCDB represent approximately 70% of newly diagnosed cases across the country. We applied for and were granted access to use the NCDB 2013 Participant User File (PUF) that was compiled on January 26, 2016. The NCDB PUF is a de-identified dataset containing cancer registry data collected from the CoC member sites.

"VUMC is a tertiary care academic medical center with 906 beds for general medical and surgical purposes, and approximately 49,000 admissions, 22,000 inpatient surgeries, 30,500 outpatient surgeries, and 108,000 emergency room visits annually. The cancer registry at VICC/VUMC is a manually curated, structured source of data on cancer patients' demographics, diagnosis, treatments, and survival outcomes. Based on CoC guidelines, patients are required to be entered into the cancer registry when their "class of case" indicates that at least the initial diagnoses or all or part of the first course of therapy are conducted at the home institution. Reportable cases must be entered into the system no later than six months after they are deemed eligible. The cancer registry at VICC has been certified by the Commission on Cancer, which designates the system's high performance in case identification and annotation. The cancer registry is a highly curated and structured data source that represents a subset of the cancer patients seen at VUMC" (Atreya, 2015).

Cancer registries are encoded using the NAACCR cancer registry data dictionary, the Facility Oncology Registry Data Standards (FORDS) manual from the Commission on Cancer (CoC), and the AJCC staging standards to codify information on cancer patients' demographics, tumor characteristics, and treatment information. At Vanderbilt, a team of specially trained nurse registrars maintains the database using the METRIQ cancer registry data management

system from Elekta.

"In addition to the data from the tumor registry system, we also leveraged clinical and administrative data from the VUMC Research Derivative (RD) (Danciu et al., 2014). The RD is a database of clinical and related data derived from VUMC clinical information systems, restructured for research, and stored on a Netezza ("IBM Netezza Data Warehouse Appliances – The Simple Data Warehouse Appliance for Serious Analytics," 2014) system. The medical record number and other identifiers are preserved within the database. Data types include reimbursement codes, clinical notes and documentation, nursing records, medication data, laboratory data, encounter and visit data, among others. Output may include structured data points, such as ICD-9-CM (International Classification of Disease) codes, CPT (Current Procedural Terminology) codes, encounter dates, semi-structured data such as laboratory tests and results, or unstructured data such as physician progress reports" (Atreya, 2015).

4.2: The Cancer Pathfinder Framework for Pathway Exploration

The general abstraction methodology for pathway exploration consists of five major subtasks: data extraction, data standardization, vertical abstraction, horizontal abstraction, and pathway exploration (Figure 5). Data extraction involves the collection of the raw, time-stamped diagnostic and treatment event data. The data standardization subtask involves structuring and mapping the information from the data extraction subtask to a set of one or more formal ontologies. The vertical abstraction subtask involves utilizing the hierarchies of the ontologies in use to identify higher-level abstractions for concepts representing the treatment events. The horizontal abstraction subtask uses clinical and temporal patterns to consolidate events over time to simplify the treatment event sequences. The pathway exploration subtask

involves allowing users to dynamically query different care paths using patient feature filtering, vertical and horizontal abstraction, and treatment filtering. This series of five subtasks converts raw data to a form that can be used to evaluate quality and support shared decision-making.

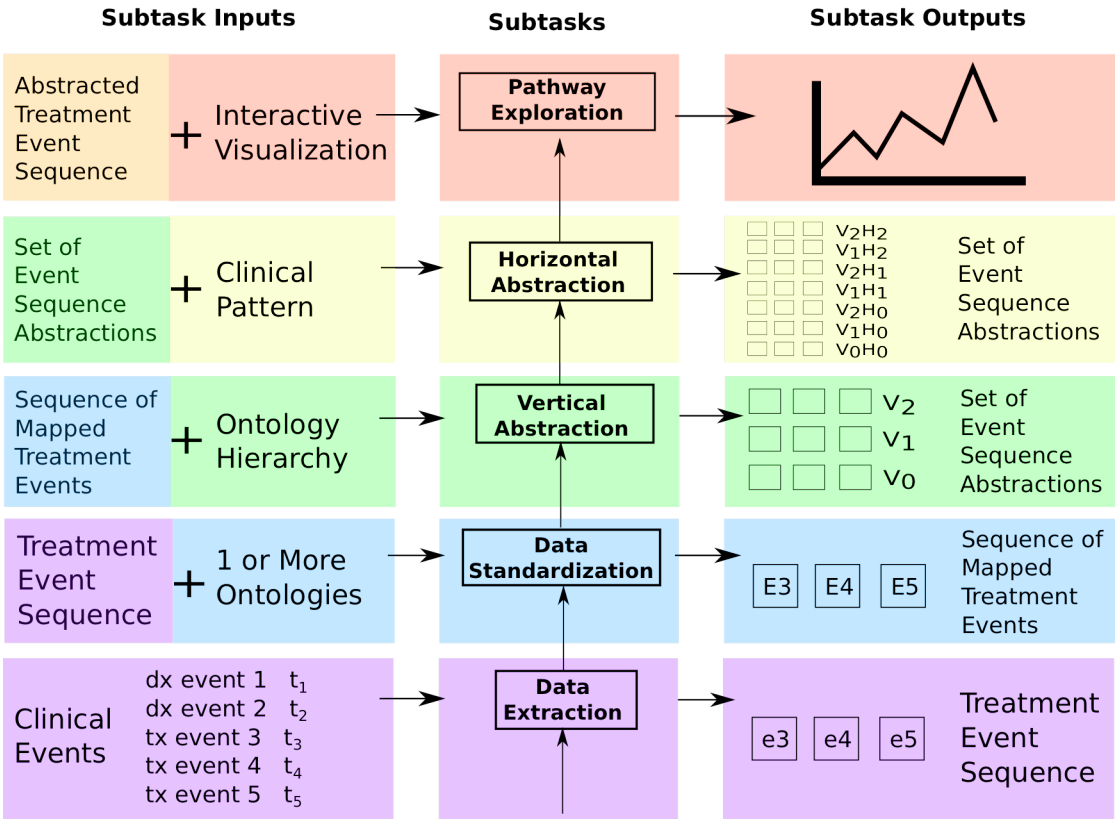

Figure 5: The five subtasks of the cancer pathfinder pathway exploration framework
The generalizable abstraction methodology utilizes five subtasks (middle) as well as their inputs (left) and outputs (right) to develop abstracted treatment event sequences from raw event data for pathway exploration. The method involves six subtasks: data extraction, data standardization, vertical abstraction, horizontal abstraction, and pathway exploration.

This abstraction methodology was implemented in the form of a web application that will be publicly available online. The raw data to be read for the data extraction subtask was pulled from the NCDB PUF File and the VUMC METRIQ cancer registry system. The event data was processed in the standardization subtask then loaded into an in-memory SQL database

(MemSQL) that was running on a non-critical server operated by VUMC IT. The vertical and horizontal abstraction subtasks are implemented in the web application in JavaScript. The interactive pathway exploration and visualization tool is implemented in JavaScript using the React.js framework and d3.js visualization library. A Node.js server uses a REST API to run queries against the MemSQL database and return results to the web application client.

4.3: Extraction and Representation of Cancer Registry Data

The data extraction subtask consisted of selecting patient records for adult women diagnosed with breast cancer from the Vanderbilt cancer registry and NCDB PUF. Diagnosis and treatment event data on biopsy, chemotherapy, immunotherapy, surgery, hormone therapy, and radiation therapy were extracted from the VUMC cancer registry and NCDB PUF. The event data in the VUMC cancer registry is more detailed than the NCDB as it provides more detail on the specific types of treatment as well as providing multiple event dates per event type. The NCDB on the other hand only provides detailed treatment information on surgical events and only contains information on a maximum of two surgical events and one of every other kind of diagnostic or treatment event. Thus, the NCDB PUF dataset is itself, already partially vertically and horizontally abstracted compared to the VUMC tumor registry. For each extracted event, the number of days since diagnosis was also calculated. We characterized the patient cohort by the patients' age at diagnosis, race, primary cancer site, histology, and clinical stage.

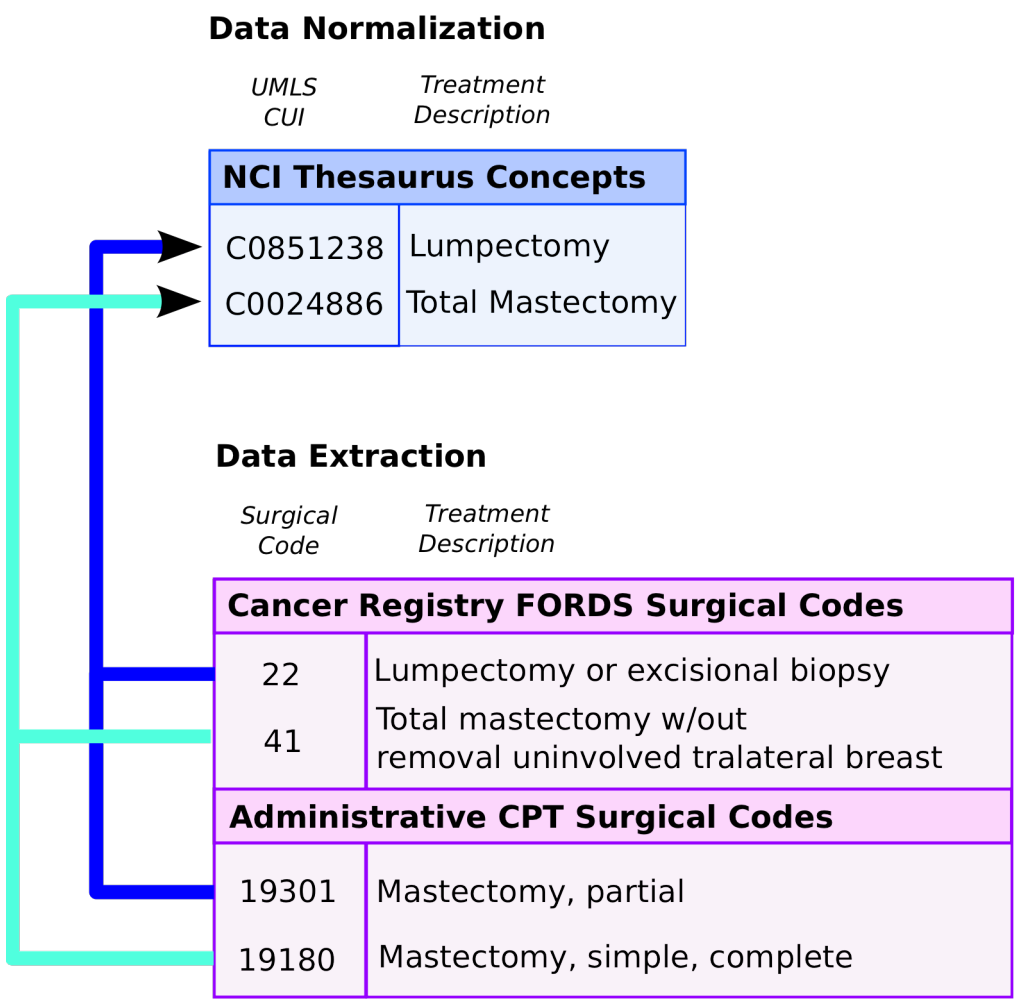

Figure 6: Example of the data extraction and normalization subtasks

Examples of FORDS manual surgical codes (i.e., 22) and CPT codes (i.e., 19301) extracted from the cancer registry and administrative data, respectively, are shown in the data extraction step. A manually derived mapping to NCI Thesaurus concepts (i.e., C15755) was used to implement the data normalization step.

"The data standardization subtask involved the mapping of treatment events from the cancer registry and administrative data to the NCI Thesaurus in order to maintain a common, formal ontology (Figure 6). The NCI Thesaurus is a specialized ontology for cancer and has hierarchies that focus on cancer specific procedures, unlike other ontologies such as SNOMED-

CT. The linkage between the cancer registry procedure terms with NCI thesaurus concepts was conducted manually by an MD/PhD student and overseen by a medical oncologist. This was necessary because linkages between the NCI Thesaurus and both CPT and FORDS codes are not publically available due to copyright restrictions from the American Medical Association (AMA) and CoC, respectively. The conversion to the NCI thesaurus allows for a more formalized abstraction process and for comparison between different data sources" (Atreya, 2015).

"The vertical abstraction subtask involved the parsing of the ontological hierarchy as well as consolidating similar events that occur simultaneously (Figure 7). A subset of the NCI Thesaurus hierarchy generated the vertical abstractions for surgical treatment events. We used three levels of vertical abstraction to represent the specific surgical terms. Base level V0 the most granular level, represents the specific type of surgical procedure such as a lumpectomy or subcutaneous mastectomy. Level V1, the next level of abstraction, represents the event subtypes Mastectomy and Breast Conservation Treatment. Level V2, the final generalization, consolidates both terms to the concept "Surgery". In addition to the surgical term hierarchy, similar treatment events occurring on the same day were consolidated to a single event. For example, two chemotherapy events listed on the same day for two separate medications (i.e., doxorubicin and cyclophosphamide) would be vertically abstracted into one chemotherapy event. The vertical abstractions are stored for use by the quality metric querying and visualization subtasks" (Atreya, 2015).

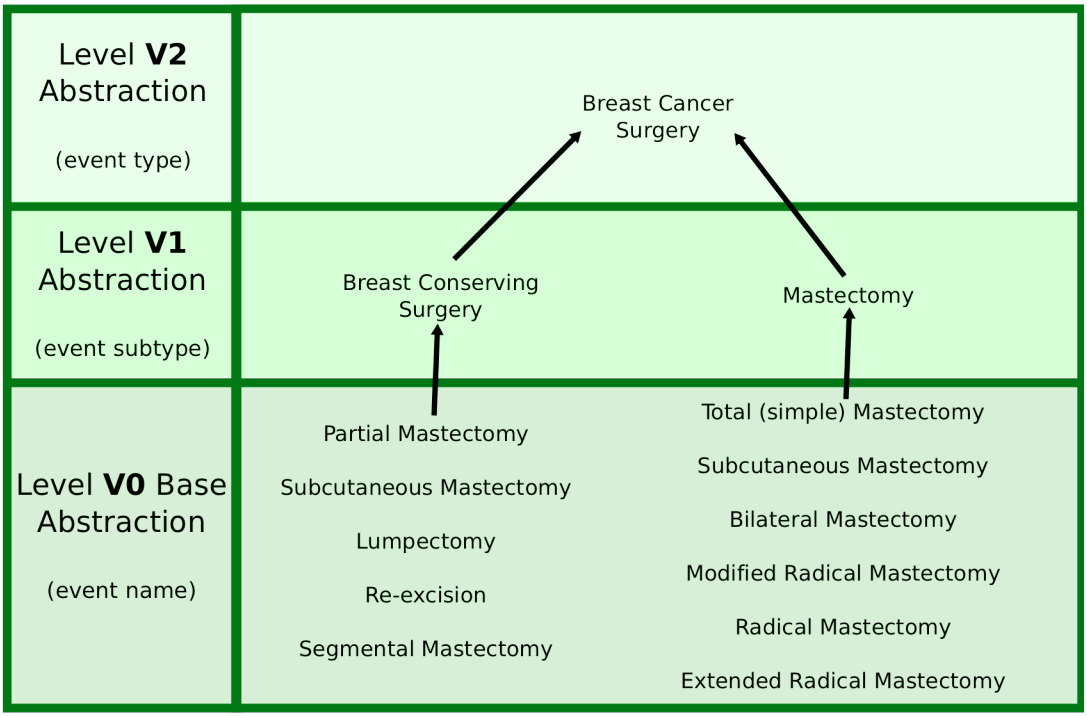

Figure 7: Example of vertical abstraction subtask for treatment event sequences

The NCI Thesaurus ontology tree for breast cancer therapeutic procedure was used for the vertical abstraction subtask to aggregate surgical events to their subtype. There are three vertical levels of abstraction represented in this hierarchy with the top tier (V2) representing a surgical procedure, the middle tier (V1) representing mastectomy versus breast conserving surgery, and the last, most granular tier (V0) representing the specific type of surgery.

"The horizontal abstraction subtask simplifies the representation of the sequence of treatment events (Figure 8). Base level H0, shows the sequence of treatment events as they occurred. The first level of abstraction, level H1, removes sequentially repeating treatment events. This allows for the representation of the sequence of treatment event types. For example, a single surgical event could represent a series of re-excision surgeries. Level H2, the final level of horizontal abstraction, represents the order of the first time any treatment event takes place. This is valuable in cancer care where the terms adjuvant (after the primary therapy) and neoadjuvant (before the primary therapy) refer to the sequence of chemotherapy with respect to surgery. As with the vertical abstractions, all horizontal abstractions are stored for quality metric querying and visualization" (Atreya, 2015).

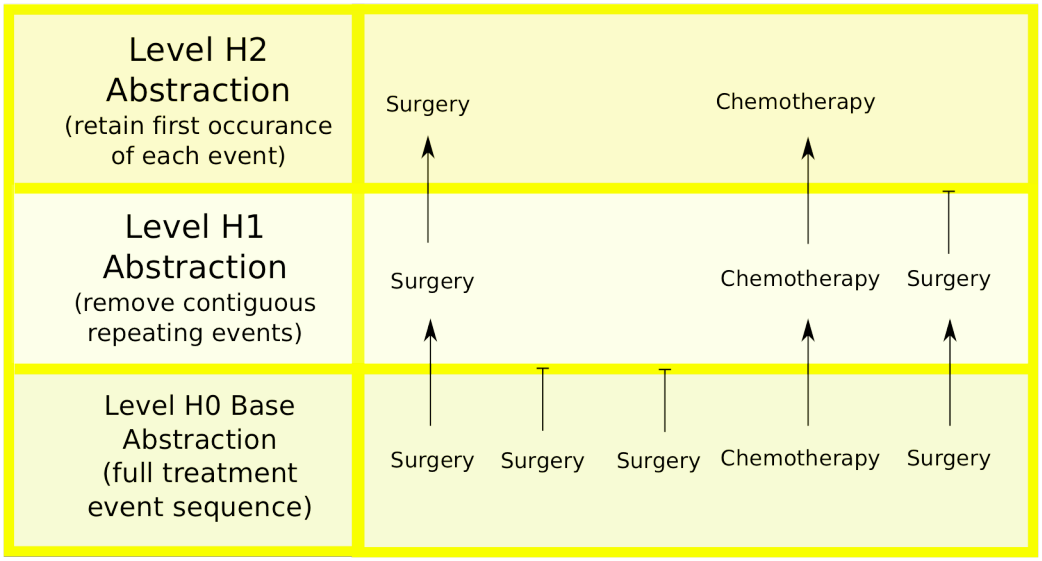

Figure 8: Example of horizontal abstraction subtask for treatment event sequences

The horizontal abstraction process represents treatment ordering on three levels. The original treatment order (H0) represents every treatment event in order at any level of vertical abstraction. The first level of abstraction (H1) removes any immediately adjacent repeating treatment events. The final level of abstraction (H2) represents only the first type of treatment event in the sequence.

After conducting the first four subtasks of our methodology, we sought to characterize the treatment event sequences for stage I-III breast cancer patients by evaluating the abstraction subtasks, visualizing the treatment paths, and searching for common treatment patterns. To evaluate abstraction process we assessed the reduction in complexity of the treatment event sequences by using three metrics: 1) we counted the number of unique treatment event sequences at each abstraction level (V0H0, V1H0, V1H1, V1H2, V2H0, V2H1, V2H2); 2) we calculated the number of treatment events per event sequence at the highest (V2H2) and lowest (V0H0) levels of abstraction to understand how the method reduces complexity; and 3) we measured how the horizontal abstraction subtask consolidated the number of treatment events from the V2H0 to V2H1 and V2H2 abstractions.

4.4: Interactive Pathway Exploration and Visualization

The pathway exploration subtask consisted of selecting Sankey visualization technique is an effective technique for representing the magnitude of flow between sequential nodes that represent diagnostic and treatment events. We represented each treatment event as a node in the Sankey diagram with the height of the node representing the relative number of patients experiencing that event. The Sankey diagram represents the sequence events but not the precise temporal duration. This visualization was implemented as part of an online web application.

We developed an interactive Sankey diagram to allow users to dynamically explore the care paths using vertical and horizontal abstractions as well as selecting specific nodes and links. Nodes represent treatment events and displayed as vertical rectangles where the height is proportional to the number of patients undergoing the treatment. Links represent the flow of

patients from node to node (event to event) where the vertical thickness of the link is proportional to the number of patients. We have implemented a variety of features to allow users to dynamically explore the Sankey diagram including feature filtering, vertical abstraction, horizontal abstraction, and treatment filtering.

We use a series of bar charts and histograms to allow users to view and filter various patient demographics, tumor and disease characteristics, and, for NCDB data, hospital characteristics. Users can dynamically interact with these features to explore the interplay of how different features interact (i.e., how does the cancer stage affect the care path). Users can vertically abstract treatment events by double-clicking on a specific node to split it into more granular events. We use the vertical abstraction hierarchies to define how the treatment event node should be split into its subcomponents. This transition will be animated for the user to see the abstraction process. Additionally, users can toggle between different levels of horizontal abstraction for VUMC registry data. The user can simplify the Sankey diagram by shortening and simplifying the event sequences in case the clinical question they are investigating does not require the full sequence complexity. Finally, users can single-click on nodes and links in the Sankey diagram to filter the population to the set of patients that had an event in a specific position (i.e., mastectomy as the second treatment event) or event transition in a specific position (i.e., breast conserving surgery as the first treatment event to mastectomy as the second).

To support the pathway exploration tool, care path events and other patient features were loaded in a MemSQL database. Care path events and event transitions were loaded at the V1H0 level of abstraction in sequence. For example, the first, second, and subsequent treatment events were stored in separate columns. Additionally, the first and subsequent event transitions (treatment event one to event two) were also stored in separate columns. This

enables users to filter on specific events and transitions by filtering the table.

4.5: Pathway-Based Quality Measurement From Cancer Registry Data

The interactive Sankey diagram tool enables users to filter on patient characteristics, treatment type, and treatment path abstraction levels. The user can navigate to the necessary vertical and horizontal level of abstraction needed to answer a specific clinical question. Users are able to assess a measure of quality from the interactive pathway exploration tool using either a basic visual assessment or a specific comparative tool.

A simple visual assessment can help identify the frequency of certain care paths in a Sankey diagram. Users could filter based on a specific type of treatment for a specific patient population and look at the relative frequency of the resulting care paths. For example, to assess the re-excision rate quality measure, a user could filter the Sankey diagram to show only surgical events, double-click on the event nodes to show the type of surgery in more detail, and then single-click on the first breast conserving surgery node in the sequence to select the patients that began their surgical course of care with a breast conserving surgery.

We will characterize the effectiveness of our pathway-based quality measurement by evaluating its ability to represent breast cancer quality measures that have been published using NCDB data and that are used by accreditation organizations such as QOPI, NQMBC, and RQRS. We will aim to replicate the analysis conducted in breast cancer quality studies using the NCDB. Many of these studies are listed in the NCDB bibliography and we have also selected additional recent publications. We will also lead a discussion section at a breast cancer tumor board meeting to get feedback from healthcare providers.

4.6: Data-Driven Decision Aids

We have developed a web-based educational tool to inform patients about their treatment options by displaying information based on the treatment pathways of previous breast cancer patients at Vanderbilt. The educational intervention is a multi-page website that briefly introduces the patient to topics, terminology, and information related to breast cancer treatment. In addition to explanations, Sankey diagrams accompany the text and illustrate the pathways through treatment undergone by VUMC patients. We used cancer registry and administrative CPT codes for women diagnosed with stage I-III breast cancer between 2010-2014 at VUMC. We have designed a web-based "storyboard" to educate new breast cancer patients about their treatment options. We also reviewed our decision aid with breast cancer patient advocates to receive their feedback.

The goal of the tool is for patients to be introduced to treatment concepts and pathways so that they can be as informed as possible as they engage their providers in the shared decision making process. Moreover, by displaying the relative rates of these decisions made by a large cohort of recent (2010-2014) patients at VUMC, these visualizations allow for patients to feel supported and comfortable with their range of options knowing other women have undergone the same decisions and treatments. The multiple pages of the site cover the following topics: introduction to the types of breast cancer treatment, introduction to how the tool works, first course of treatment, breast conserving surgery (“lumpectomy”), mastectomy and reconstruction, pre-surgical drug therapy, and post-surgical drug therapy.

This data-driven decision aid is implemented in the form of a web application that will be available online. The interactive pathway exploration and visualization tool is implemented

in JavaScript using the React.js framework and d3.js visualization library. A Node.js server running on a non-critical server operated by VUMC IT is used to serve the web application.

Chapter 5: Results of the Pathfinder Methodology

We extracted breast cancer patient data from the VUMC and NCDB cancer registries (Section 5.1). We applied the abstraction method to standardize treatment event representation and then developed and evaluated the treatment event sequences (Section 5.2). We present our pathway exploration tool and describe its various features in Section 5.3. Our pathway exploration framework is robust and can represent the majority of quality measures from accreditation agencies and can replicate most of the studies conducted on NCDB data sources (Section 5.4). Finally, we describe our data-driven decision aid, which introduces newly diagnosed breast cancer patients to potential treatment options (Section 5.5).

5.1: Data Extraction from Cancer Registry for Breast Cancer Patient Cohort

We extracted the records of 1528 stage I-III female breast cancer patients diagnosed between 2000-2012 from the Vanderbilt cancer registry that had undergone some diagnostic or treatment event, had a full NAACCR record in the data warehouse, and met our inclusion criteria. The 72 experimental treatment events were mapped to their respective 68 chemotherapy events and 4 hormone therapy events. The demographic characteristics for this population are presented in Table 1.

Table 1: Demographic characteristics of the VUMC breast cancer patient cohort

Characteristic	Median	1st Quartile	3rd Quartile	Min	Max
Age at diagnosis (N=1528)	54	46	64	21	99

Characteristic	Category	Count
Race (top 3)	**Total**	**1528**
	White	1311 (85.8%)
	Black	171 (11.2%)
	Other Asian, including Asian or Oriental, NOS	16 (1.0%)
Primary Site (all)	**Total**	**1528**
	C500, Nipple	10 (0.7%)
	C501, Central portion of breast (subareolar)	85 (5.6%)
	C502, Upper inner quadrant	159 (10.4%)
	C503, Lower inner quadrant	107 (7.0%)
	C504, Upper outer quadrant	621 (40.6%)
	C505, Lower outer quadrant	120 (7.9%)
	C506, Axillary tail	4 (0.3%)
	C508, Overlapping lesion of breast	314 (20.5%)
	C509, Not otherwise specified (NOS)	108 (7.1%)
Histology (top 5)	**Total**	**1528**
	8500/3, Infiltrating duct carcinoma, NOS	888 (58.1%)
	8010/3, Carcinoma, NOS	243 (15.9%)
	8520/3, Lobular carcinoma, NOS	113 (7.4%)
	8522/3, Infiltrating duct and lobular carcinoma	76 (5.0%)
	8211/3, Tubular adenocarcinoma	63 (4.1%)
Clinical stage (all)	**Total**	**1528**
	1	854 (55.9%)
	2	503 (32.9%)
	3	171 (11.2%)

Demographic characteristics are provided for the 1528 patient records in the VUMC breast cancer patient cohort. Characteristics derived from the cancer registry include age, race, primary tumor site, histology, and clinical stage.

We extracted the records of 2032209 stage 0-IV male and female breast cancer patients diagnosed between 2004-2013 from the NCDB registry. There were 2013590 (99%) female patients in this dataset. The demographic characteristics for this population are presented in Table 2.

Table 2: Demographic characteristics of the NCDB breast cancer patient cohort

Characteristic	Median	1st Quartile	3rd Quartile	Min	Max
Age at diagnosis (N=2032209)	60	51	71	18	90

Characteristic	Category	Count
Race (top 3)	**Total**	**2032209**
	White	1698317 (83.6%)
	Black	229019 (11.3%)
	Unknown	23787 (1.2%)
Primary Site (all)	**Total**	**2032209**
	C500, Nipple	13693 (0.7%)
	C501, Central portion of breast (subareolar)	111222 (5.5%)
	C502, Upper inner quadrant	208432 (10.3%)
	C503, Lower inner quadrant	113093 (5.6%)
	C504, Upper outer quadrant	672861 (33.1%)
	C505, Lower outer quadrant	139850 (6.9%)
	C506, Axillary tail	8445 (0.4%)
	C508, Overlapping lesion of breast	432734 (21.3%)
	C509, Not otherwise specified (NOS)	331879 (16.3%)
Histology (top 5)	**Total**	**2032209**
	8500/3, Infiltrating duct carcinoma, NOS	1185587 (58.3%)
	8500/3, Intraductal carcinoma, noninfiltrating NOS	160411 (7.9%)
	8520/3, Lobular carcinoma, NOS	143317 (7.1%)
	8522/3, Infiltrating duct and lobular carcinoma	88533 (4.4%)
	8523/2 Infiltr. duct mixed with other types of carcinoma, in situ	76694 (3.8%)
Clinical stage (all)	**Total**	**2032209**
	0	334413 (16.5%)
	1	312954 (15.4%)
	2	336917 (16.6%)
	3	96144 (4.7%)
	4	71620 (3.5%)

Demographic characteristics are provided for the 2,032,209 patient records in our NCDB breast cancer patient cohort. Characteristics derived from the cancer registry include age, race, primary tumor site, histology, and clinical stage.

We compared the number of treatment events in the VUMC and NCDB datasets for stage I-III patients (Table 3). The 1528 VUMC registry patient cases had 4893 treatment events including 1812 surgical events, 1239 chemotherapy events, 958 hormone therapy events, and 884 radiation therapy events with an average of 3.2 treatment events per patient. These 1474032 NCDB cases had 3878189treatment events including 1691571 definitive surgical events, 589236 chemotherapy events, 795575 hormone therapy events, and 801807 radiation therapy events with an average of 2.6 treatment events per patient. This difference in the number of treatment events per patient is due to the pre-horizontally abstracted nature of the

NCDB treatment event data. The pre-abstracted nature of the NCDB data leads to fewer treatment events per patient which is highlighted by the 0.8 chemotherapy events per patient present in the VUMC registry compared to 0.4 in the NCDB.

Table 3: Number of treatment events in NCDB and VUMC registries

	NCDB	VUMC Registry
# of Patients	1474032	1528
Treatment Events	**# events (# events per patient)**	
Total # of Events	3878189 (2.6)	4893 (3.2)
Surgical events	1691571 (1.1)	1812 (1.2)
Chemotherapy	589236 (0.4)	1239 (0.8)
Hormone therapy	795575 (0.5)	958 (0.6)
Radiation therapy	801807 (0.5)	884 (0.6)

The number of treatment events and treatment events per patient present in the NCDB and VUMC cancer registry for stage I-III patients. There are 2.6 events per patient in the NCDB compared to 3.2 in the VUMC registry. The greatest difference is for the 0.4 chemotherapy events in the NCDB compared to the 0.8 chemotherapy events per patient in the VUMC registry.

5.2: Data Normalization and Abstraction from Cancer Registry Data

The standardization process simplified the representation of surgical events from 35 FORDS cancer registry surgical codes to 8 NCI Thesaurus clinical concepts for the primary surgical treatments in the NCDB and VUMC registry. The mapping between the FORDS and NCI thesaurus concepts can be found in Supplemental Table 1.

The lowest level (V_0H_0) of abstraction yielded 391 unique treatment event sequences while the highest level (V_2H_2) of abstraction had 45. The vertical and horizontal abstraction process provided an 8.7-fold reduction in the number of unique treatment strings. a 42% decrease (3.1 bits) in Shannon entropy, and 8.3-fold decrease in diversity from the least abstracted (V_0H_0) to most abstracted (V_2H_2) representation (Table 4). The analysis of Shannon's

entropy shows how the amount of information contained in an event sequence is reduced from 8.3 bits at abstraction level V0H0 to 4.2 bits at level V2H2. The diversity, which is derived from the Shannon entropy, shows how the effective number of treatment event sequences is reduced from 156.3 event sequences at the V0H0 abstraction level to 18.9 at the V2H2 level. The greatest percent decrease from actual to effective number of sequences between abstraction levels V0H0 and V2H2 occurred at the V2H1 level of abstraction while the smallest occurred at the V1H2 level. The fraction of event sequences needed to cover 90% of the population decreased from 0.63 to 0.36. We also counted the number of most frequently occurring treatment event sequences that represented 90% of the patient population. We showed how the fraction of event sequences needed to cover 90% of the population decreased from 62.7% at V0H0 to 35.5% at V2H2. The greatest relative concentration of sequences occurred at the V_2H_1 level of abstraction while the lowest occurred at the V_1H_2 level. The distribution of treatment event sequence frequency exhibited an exponential distribution with a long tail that was compacted through the abstraction process (Figure 9).

Table 4: Treatment event sequences in VUMC registry

Level of abstraction	# of unique treatment event sequences	# (%) of unique treatment event sequences to cover 90% of population	Shannon entropy for set of unique treatment event sequences (bits per event sequence)	Diversity of event sequences based on Shannon's entropy (effective # of sequences)	% decrease from # of actual to effective treatment event sequences
Raw data	632	480 (75.8%)	8.3	322.7	48.9%
Level V_0H_0 abstraction	391	245 (62.7%)	7.3	156.3	60.0%
Level V_1H_0 abstraction	265	126 (47.5%)	6.4	86.6	67.3%
Level V_2H_0 abstraction	205	79 (38.5%)	5.9	60.1	70.7%
Level V_1H_1 abstraction	154	54 (35.1%)	5.5	45.5	70.5%
Level V_2H_1 abstraction	97	26 (26.8%)	4.7	26.0	73.2%
Level V_1H_2 abstraction	73	28 (38.4%)	4.9	30.3	58.5%
Level V_2H_2 abstraction	45	16 (35.5%)	4.2	18.9	58.0%

The abstracted treatment event sequences for 1528 stage I-III VUMC breast cancer patients showed a 8.7-fold reduction in the number of unique treatment strings, a 42% decrease in Shannon entropy, and 8.3-fold decrease in effective diversity from the least abstracted (V0H0) to most abstracted (V2H2) representation.

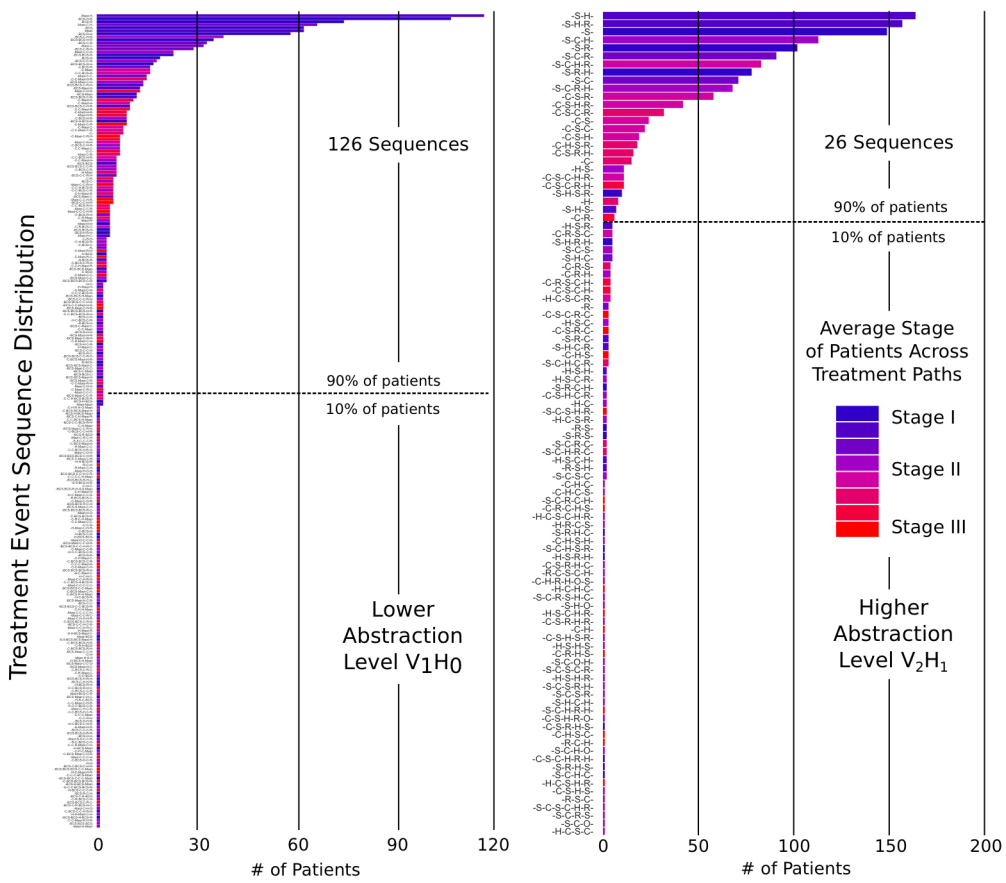

Figure 9: Distribution of treatment event sequences across patients at abstraction levels V1H0 and V2H1

The number of VUMC patients per treatment event sequence at two levels of abstraction (V1H0 on left, V2H1 on right) is shown in descending order (top to bottom) and demonstrates an exponential distribution. The color of the bar represents the average cancer stage of the patients with a specific treatment path (red represents a higher stage compared to blue). The lower level of abstraction (V1H0) has 265 event sequences in which 47.5% of the sequences are required to represent 90% of the patient population. The higher level of abstraction (V2H1) has 97 event sequences in which 26.8% of sequences are required to represent 90% of the patient population.

The NCDB provides pre-abstracted treatment data where either one or two surgical events and only one event of every other kind of treatment are provided. The vertical and part-horizontal abstraction process a 5.6-fold reduction in the number of unique treatment strings

from 353 to 65 unique sequences. The abstraction process also led to a 27% decrease in Shannon entropy, and a 2.6-fold decrease in diversity from the least abstracted (V1H1*) to the most abstracted (V2H2) representation (Table 5). The diversity, which is derived from the Shannon entropy, shows how the effective number of treatment event sequences is reduced from 38.7 event sequences at the V1H1* abstraction level to 14.8 at the V2H2 level. The greatest percent decrease from actual to effective number of sequences between abstraction levels V1H1* while the smallest occurred at the V2H2 level. The greatest relative concentration of sequences occurred at the V1H1* level of abstraction while the lowest occurred at the V2H2 level. The distribution of treatment event sequence frequency had an exponential distribution (Figure 10). We showed how the fraction of event sequences needed to cover 90% of the population increased from 10.8% at V0H0 to 16.9% at V2H2. The greatest relative concentration of sequences occurred at the V2H1 level of abstraction while the lowest occurred at the V1H2 level.

Table 5: Treatment event sequences in NCDB registry

Level of NCDB surgical abstraction	# of unique treatment event sequences	# (%) of unique treatment event sequences to cover 90% of population	Shannon entropy for set of unique treatment event sequences (bits per event sequence)	Diversity of event sequences based on Shannon's entropy (effective # of sequences)	% decrease from # of actual to effective treatment event sequences
Level V_1H_{1*} abstraction	353	38 (10.8%)	5.3	38.7	89.0%
Level V_1H_2 abstraction	149	21 (14.1%)	4.5	23.4	84.3%
Level V_2H_{1*} abstraction	163	22 (13.5%)	4.6	24.5	85.0%
Level V_2H_2 abstraction	65	11 (16.9%)	3.9	14.8	77.2%

The abstracted treatment event sequences for 1474032 stage I-III NCDB breast cancer patients showed a 5.6-fold reduction in the number of unique treatment strings, a 27% decrease in Shannon entropy, and 2.6-fold decrease in effective diversity from the least abstracted (V1H1*) to most abstracted (V2H2) representation.

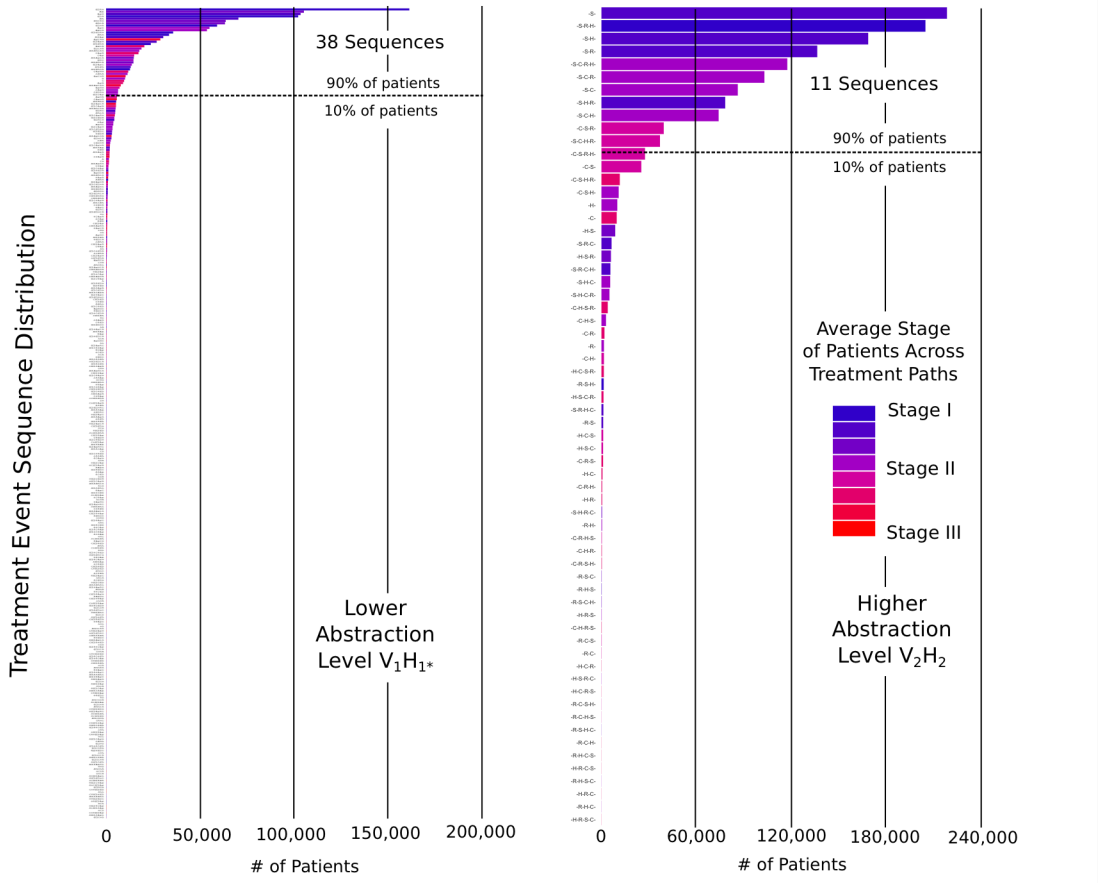

Figure 10: Distribution of treatment event sequences across patients at abstraction levels V1H1* and V2H2

The number of NCDB patients per treatment event sequence at two levels of abstraction (V1H1* on left, V2H2 on right) is shown in descending order (top to bottom) and demonstrates an exponential distribution. The color of the bar represents the average cancer stage of the patients with a specific treatment path (red represents a higher stage compared to blue). The lower level of abstraction (V1H1*) has 365 event sequences in which 9.9% of the sequences are required to represent 90% of the patient population. The higher level of abstraction (V2H2) has 65 event sequences in which 18.5% of sequences are required to represent 90% of the patient population.

Many of the top 10 treatment event sequences in the NCDB and VUMC registry are similar at the V2H2 level of abstraction (Table 6). Of the top 10 sequences, all but one in each list are present in the other. The only exceptions in each list (S-C-R-H in the NCDB and S-C-H-R in the NCDB) are present in the 11th spot in the other. Of the top 10 sequences in each list, only

the sequence C-S-R in the VUMC registry is the only sequence involving neo-adjuvant chemotherapy while all other sequences began with a surgical treatment.

Table 6: Top 10 treatment event sequences in the NCDB and VUMC registries

NCDB Top 10 Sequences		VUMC Top 10 Sequences	
Sequence	**Count (%)**	**Sequence**	**Count (%)**
S	218886 (15.3%)	**S-H-R**	172 (11.3%)
S-R-H	205252 (14.4%)	**S-H**	171 (11.2%)
S-H	168976 (11.8%)	**S**	149 (9.8%)
S-R	136714 (9.6%)	**S-C-H**	114 (7.5%)
S-C-R-H	117812 (8.2%)	**S-R**	104 (6.8%)
S-C-R	103277 (7.2%)	**C-S-R**	96 (6.3%)
S-C	86396 (6.0%)	**S-C-R**	96 (6.3%)
S-H-R	78505 (5.5%)	**S-C-H-R**	93 (6.1%)
S-C-H	74262 (5.2%)	**S-R-H**	79 (5.2%)
C-S-R	39568 (2.8%)	**S-C**	79 (5.2%)

The top 10 treatment event sequences for stage I-III patients in the NCDB and VUMC registries are listed at the V2H2 abstraction level along with the count and percentage for each sequence.

5.3: Pathway Exploration

The interactive pathway exploration tools for NCDB and VUMC cancer registry data utilize Sankey diagrams and other charts representing patient, tumor, and hospital features. This pathway exploration tool can be accessed at https://www.cancerpathfinder.org. Each Sankey and chart element is filterable allowing users to drill into the underlying data. Our tool enables a variety of features including vertical abstraction, horizontal abstraction, treatment filtering, patient and tumor features, and hospital features. We will walk through various important patient features to demonstrate the value of the tool features (Table 7).

Table 7: Treatment path filtering examples

Filters	Dataset	Fraction of First Treatment in Path			
		Surgery	Chemotherapy	Hormone therapy	Radiation therapy
No filters	NCDB	89.8%	6.9%	2.5%	0.8%
HER2 positive	NCDB	78.0%	18.8%	2.1%	1.0%
Clinical & pathologic stage I	VUMC	87.8%	2.6%	4.9%	4.6%
Clinical & pathologic stage III	NCDB	65.3%	31.1%	3.2%	0.3%
Clinical & pathologic stage III and triple negative	NCDB	60.4%	38.8%	0.4%	0.4%
Clinical & pathologic stage I	VUMC	65.4% BCS 20.6% Mastectomy 1.8% Unknown surgery	2.6%	4.9%	4.6%
First event only & clinical + pathological stage I	VUMC	65.4% BCS 20.6% Mastectomy 1.8% Unknown surgery	2.6%	4.9%	4.6%

The distribution of the first step of the treatment paths described in this section are displayed.

Users can filter on the patient population using a variety of patient and tumor characteristics that are represented by a set of bar charts and histograms. In the NCDB dataset, users will also be able to filter on hospital characteristics. To demonstrate this feature, we walk through a series of examples. We first look at how focusing on HER2 positive patients affects the distribution for the year of diagnosis as most testing has been done since 2010. We next compared the treatment paths and the fraction of patients that undergo neo-adjuvant chemotherapy for stage I and stage III patients. A larger fraction of patients with stage III cancer undergo neo-adjuvant chemotherapy in order to shrink the tumor before a potential surgery. This fraction grows for patients with stage III, triple negative breast cancer. We demonstrated vertical abstraction by splitting the surgical concept into breast conserving surgery and mastectomy to demonstrate their divergent treatment path.

We have enabled dynamic vertical abstraction to aid in the exploration of VUMC and

NCDB cancer registry. The tool defaults to showing all treatment events at the highest (V2) level of abstraction. If there is a lower level of abstraction available, the user is able to double click on the vertical bar of a Sankey diagram that represents an event. The bar then splits into multiple vertical bars representing the respective next lower level of vertical abstraction (V1). This next level would show separate vertical treatment bars for breast conserving surgery (BCS) and mastectomy. For example, stage I breast cancer patients who begin with surgical treatment undergo a BCS 65% and a mastectomy 21% of the time.

Users can also filter out certain kinds of treatment events from the Sankey diagram to help reduce the amount of information being presented and focus on the interaction of specific modalities of treatment. For example, a user could choose to view only surgical paths if they had a surgery specific question (i.e., re-excision rate). Additionally, if a user wants to identify the rate of radiation therapy after lumpectomy, the user could filter out hormone therapy and chemotherapy.

Users are able to dynamically alter the horizontal abstraction level by using a set of toggle buttons near the Sankey diagram for VUMC cancer registry data. The users can choose to view every treatment event in the path, only non-repeating events, or only one event of every kind.

The pathway exploration platform was presented to a meeting of the Vanderbilt Breast Center tumor board in order to get their feedback. The interdisciplinary team of breast cancer providers believed that the tool could be very valuable support for research and quality improvement efforts of Vanderbilt and national cancer datasets. Providers also stressed that to use the tool with patients, they would need to easily access simplified treatment paths as shown in the patient education tool (see section 5.5). The version of the tool that was presented to the

tumor board was prior to the implementation of the treatment filtering feature and the feedback received demonstrated the importance of that feature in simplifying the Sankey diagram to make it easier to comprehend in a clinical setting.

5.4: Pathway Based Quality Measurement

We determined that all of the QOPI, NQMBC, and RQRS measures we assessed could be effectively represented by our abstraction method with the addition of temporal constraints and more event types (Table 9, 10). Without any additional functionality, the current framework could measure 14 of 20 measures that involve the relative frequency of treatment events, the focus of this work. The other six measures require calculation of the absolute temporal duration between events. There are 46 additional measures that utilize additional event types including diagnostic (i.e., imaging, biopsy, labs), clinical encounter (i.e., pathology report, proper assessment conducted), and outcome (i.e., complications, survival) events that were not the focus of this work. If we represented these events in our framework, we could represent 32 of the 46 measures. The remaining 14 measures would again require the implementation of specific temporal constraints.

Table 8: Quality measure representation of cancer pathfinder framework

Quality Measure Pattern Type	Framework's Ability to Represent Measures				Total
	Complete	Partial		Unable to Represent Measure	
		Other Event Types	Temporal Constraints		
Treatment Events Only	14 (70%)	0 (0%)	6 (30%)	0 (0%)	20 (30%)
Treatment & Other Events	0 (0%)	12 (100%)	7 (58%)	0 (0%)	12 (18%)
Other Events Only	0 (0%)	34 (100%)	7 (21%)	0 (0%)	34 (52%)
All Quality Measures	14 (21%)	46 (70%)	21 (32%)	0 (0%)	66 (100%)

The ability of this framework to either completely, partially, or not represent quality measures from QOPI, NQMBC, and RQRS is shown below. If the framework can partially represent the measure, then the additional feature necessary to completely represent the measure is listed.

We evaluated previously published breast cancer studies that were based on the NCDB to determine if our pathway exploration framework could help a clinical researcher answer the question. We identified 23 references in NCDB bibliography from 2006-2013 as well as 7 additional articles using NCDB data from 2013-2016. For each study, we assessed if the measure of quality or clinical question discussed in the article could be evaluated using our pathway exploration framework. If studies that could not be re-created, we identified what missing features would be necessary to include in the system. While only five of the 30 current studies can currently be replicated using the pathway exploration tool in its current state, most others are easily added by including more features from the dataset. Four of the 30 studies will also require the addition of temporal features to select specific time intervals between events. (Table 10).

Table 9: Breast cancer clinical research studies in NCDB registry

Clinical Study	Framework Ability to Represent Study	Feature Needed to Enable Full Representation
Ethnic distribution for clinical trial accrual (Newman et al., 2006)	Partial	Patient feature filter
Insurance status vs. early or advanced stage of	Partial	Patient feature filter

disease at diagnosis (M. T. Halpern, Bian, Ward, Schrag, & Chen, 2007)		
Treatment trends and survival for T1 cancers (Kennedy et al., 2007)	Partial	Patient feature filter
Treatment for metaplastic breast cancer (Pezzi et al., 2007)	Partial	Patient feature filter
Rate of sentinel lymph node biopsy (A. Y. Chen et al., 2008)	Partial	Patient feature filter
Rate of sentinel lymph node biopsy and axillary lymph node dissection (M. Halpern, Chen, Marlow, & Ward, 2009)	Partial	Patient feature filter
Race vs. tumor characteristics (Desantis, Jemal, & Ward, 2010)	Partial	Patient feature filter
Time between surgery and adjuvant chemotherapy (S. A. Fedewa, Ward, Stewart, & Edge, 2010)	Partial	Temporal constraint
Rate of contralateral prophylactic mastectomy (Yao, Stewart, Winchester, & Winchester, 2010)	Partial	Treatment event detail
Time to initial treatment post-biopsy (S. a Fedewa et al., 2011)	Partial	Temporal constraint
Race and insurance status vs. hormone receptor testing, regional therapy, and chemotherapy (Freedman et al., 2011)	Partial	Patient feature filter
Geography and institution vs. race vs. early or late stage of disease at presentation (Keller, Guilfoyle, & Sariego, 2011)	Partial	Patient feature filter
Diagnosis via needle biopsy vs. excision in T3 cancers (Williams et al., 2011)	Partial	Patient feature filter
Age vs. mastectomy or breast conserving surgery (Freedman et al., 2012)	Complete	Patient feature filter
Compare male vs. female breast cancer (Greif, Pezzi, Klimberg, Bailey, & Zuraek, 2012)	Partial	Patient feature filter
TNM stage vs. post-mastectomy radiation (Huo et al., 2012)	Complete	Patient feature filter
Age vs. post-lumpectomy radiation (K. et al., 2012)	Complete	Patient feature filter
Socioeconomic factors vs. rate of reconstructive surgery (Sisco et al., 2012)	Partial	Patient feature filter
Specific radiation therapy factors (Czechura et al., 2013)	Partial	Treatment event detail
Race vs. sentinel lymph node biopsy and axillary lymph node dissection (Black, Jiang, Kuerer, Buchholz, & Smith, 2014)	Partial	Treatment event detail
Specific radiation therapy factors (Wang et al., 2014)	Partial	Treatment event detail
Time to surgery/treatment vs. unilateral and bilateral mastectomy (Sharpe et al., 2014)	Partial	Add temporal features
Treatment vs. survival in inflammatory disease (Rueth et al., 2014)	Partial	Patient feature filter
Guideline concordance on radiation therapy for DCIS patients (Yao et al., 2014)	Partial	Patient feature filter
Tumor size and lymph node ratio - # positive / # examined lymph nodes (Wiznia et al., 2014)	Partial	Patient feature filter
Mastectomy rate in lumpectomy-eligible	Complete	Patient feature filter

population (Kummerow, Du, Penson, Shyr, & Hooks, 2015)		
Race vs. chemotherapy usage and pathologic complete response (Killelea et al., 2015)	Partial	Patient feature filter
Post-mastectomy radiation usage (Minami et al., 2016)	Complete	
Survival and time to surgery from diagnosis (Bleiche et al., 2016)	Partial	Temporal constraint
No neo-adjuvant chemotherapy with lumpectomy and regional nodal irradiation (Hou, Yao, Jaskowiak, Hasan, & Winchester, 2012)	Partial	Treatment event detail

Clinical breast cancer studies using NCDB data are listed. We determined if our pathway exploration framework can be used to answer the same clinical question evaluated in the previously published studies. We list categorized the ability to answer the question as complete (can represent the quality measure), partial (can represent the quality measure with the addition of a simple feature), and not possible (not possible to represent the quality measure). In the case of partial determination, the additional feature necessary was listed.

We replicated parts of a 2012 study that assessed post-mastectomy radiation therapy for patients with one to three positive lymph nodes. The study demonstrated that between 1998-2007, approximately 20% of T1-2, N1 patients underwent radiation therapy after mastectomy and that age was inversely correlated as 31.3% of patients under age 40 received radiation compared to 8.2% of patients above the age of 80 (Huo et al., 2012). We were able to show similar proportions of state II breast cancer patients between 2004-2007 having post-mastectomy radiation. Across all ages, 21.9% of stage II patients (4098 out of 50,293) between 2004-2007 had post mastectomy radiation (33.6%, 555 of 4,650 under the age of 40 & 8.2%, 189 of 4,880 over the age of 80) (Figure 11). In addition, we were able to update the analysis to show a slight increase in the rate of post-mastectomy radiation. Between 2010-2013, 26.1% of stage II breast cancer patients underwent post-mastectomy radiation therapy (40%, 1,475 of 7,071 patients under age 40 & 9.9%, 534 of 10,980 patients above age 80).

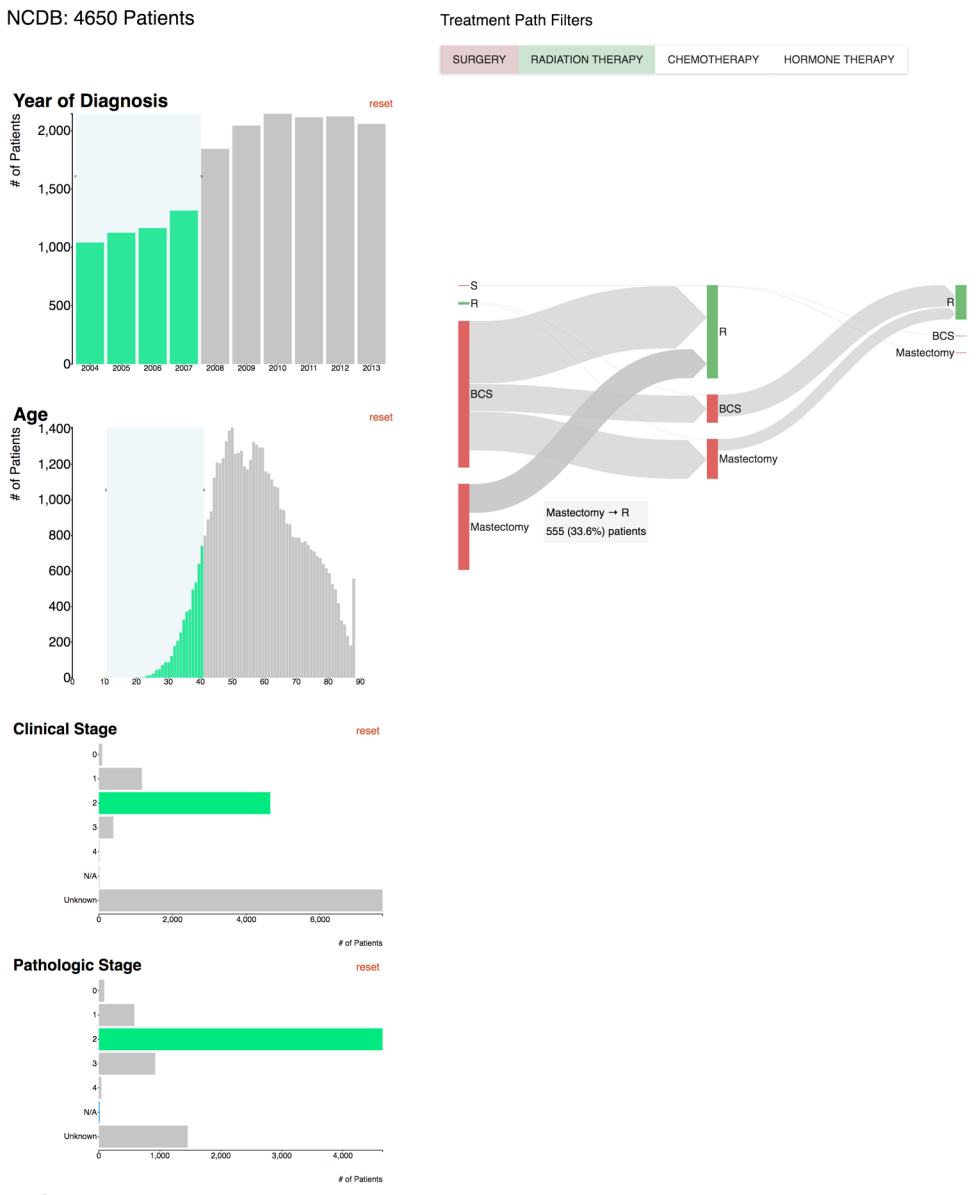

Figure 11: Partial replication of a 2012 study that assessed post-mastectomy radiation therapy.

The pathway exploration can show that fraction of stage II breast cancer patients under the age of 40 between 2004-2007 that underwent post-mastectomy radiation.

5.5: Data-Driven Decision Aids

We developed an educational, data-driven decision aid with seven web-pages that introduces new breast cancer patients to the modalities of treatment, explains how Sankey diagrams are interpreted, outlines the first course of treatment, introduces local therapy in the context of lumpectomy and mastectomy, and discusses pre- and post-surgical drug therapy. This decision aid can be accessed at https://www.mypathfinder.app.vumc.org. We used VUMC cancer registry and administrative CPT codes for women diagnosed with stage 0-III breast cancer between FY2010-14 to generate the Sankey diagrams. In order to simplify the diagrams for patients, we faded in new sections of a Sankey diagram as the patient scrolls through the screen.

The first page of the educational decision aid is focused on introducing the different modalities of treatment. The second page of the decision aid introduces the Sankey diagram concept and how it will represent various treatment options. A simple mock treatment path was developed to show how a treatment path could diverge and converge as patients complete their care at different steps. The third page presents a broad overview of how breast cancer treatment begins either with a surgery or with chemotherapy. We walk patients through two separate Sankey diagrams that begin with lumpectomy, mastectomy, or chemotherapy. For patients that begin with surgery, we also visualize subsequent drug therapies. We simplified this Sankey diagram to always show that hormone therapy begins after radiation therapy to avoid the visual complexity for few cases where they begin near the same time. The fourth page introduces the lumpectomy surgery and related local therapies. The surgical path introduces the concept of a re-excision as well as the case where the patients have a mastectomy after an

initial lumpectomy. Additionally, this page shows how a large fraction of patients that undergo a lumpectomy go on to have local radiation therapy. This section also links out to various supporting websites that provide more detailed information on the specific forms of treatment. The fifth page walks through the decision making process for the mastectomy procedure as well as breast reconstruction. The Sankey diagram walks through option of unilateral and bilateral mastectomy and the choice of having or bypassing reconstructive surgery. The Sankey diagram also shows the rates of the types of breast reconstruction. Finally, the Sankey diagram shows how a small fraction of this population has radiation therapy. The sixth page presents the concept of neo-adjuvant chemotherapy when focusing on pre-surgical drug therapy. The Sankey diagram shows the fraction of patients that begin with chemotherapy and go on to have either a lumpectomy or mastectomy. The seventh page displays post-surgical therapy as it describes adjuvant chemotherapy and hormone therapy after a lumpectomy or mastectomy.

The data-driven decision aid was presented to both a meeting of the Vanderbilt Breast Center tumor board as well as a group of breast cancer patient advocates affiliated with the Vanderbilt-Ingram Cancer Center in order to get their feedback. Both groups provided positive feedback that the tool effectively presented educational material in a story format with simple visualizations that patients could understand and read through even prior to their initial appointment. The patient advocates appreciated the ease of use of the tool as it only required scrolling down the page to see the visualizations develop. The discussion at the tumor board focused on the nature of the materials patients access on the internet prior to their appointment. While some providers were initially concerned about the patients receiving information without the provider present, others commented that the patients already were seeing information online and this resource would be well curated

Chapter 6: Conclusions and Discussion

6.1: Informatics Contributions

We have developed an abstraction framework that simplifies raw treatment event sequences to overcome challenges in directly exploring care pathways. It enables the representation of patterns of care, the evaluation of desired quality measures, and development of decision-aids for our breast cancer patient population. We mapped surgical events to clinical concepts and developed abstraction rules based on conceptual hierarchies and event sequence patterns. We were able to analyze, visualize, and explore treatment paths across our population by building a pathway exploration tool based on our abstraction methodology. We gained insight into how the abstraction process compacts treatment event sequences and provides different views of the variability across treatment paths in the population. We demonstrated how this framework could generate quality measures by identifying a reduction in re-excision rate over the past decade.

Prior work in temporal abstraction has often focused abstracting events to a specific level of granularity in order to generate a specific decision support item. By representing treatment event sequences at multiple levels of vertical and horizontal abstraction we can study the different data representations to understand how each abstraction level portrayed treatment paths. We can then dynamically query them based on the clinical question at hand.

Analysis of the abstractions for VUMC cancer registry data showed how aggregated treatment event paths could simplify the information's representation and enable effective exploration of the data. While the majority of treatment event sequences were aggregated,

some highly-variable care in the long tail of the distribution was not simplified through the abstraction process. This long tail was most compacted when represented only by the relative order of treatment events (H2 abstraction) due to the smaller relative decrease in the number of actual to effective number of sequences. On the other hand, the greatest concentration of frequently occurring treatment paths and elongation of the long tail occurred when using the full ordering of non-repeating treatment events (H1 abstraction). The analysis of the abstraction process demonstrates how the simplification of treatment pathways can lead to variable sizing and diversity of patient cohorts.

The abstraction of the NCDB treatment paths was different because the NCDB included pre-abstracted treatment events. For example, only two surgical events were provided (first and definitive) and the type of surgery was only provided at the V0 level of abstraction for the definitive one. Additionally, other treatment events were pre-abstracted at the V2H2 abstraction level. As a result, the results are not directly comparable to the VUMC assessment. The effective number of sequences provided a means for comparing the NCDB abstractions. The V2H2 abstractions between the NCDB and VUMC were similar (14.8 and 18.9, respectively). The closest match the original NCDB event sequence (V1H1*) with 38.7 effective event sequences is in between VUMC's V1H2 and V1H1 levels of abstraction (30.3 and 45.5 sequences, respectively). The V1H1* abstraction level led to the greatest consolidation of treatment sequences. The number of unique treatment event sequences outpaced the effective number of sequences when moving to higher levels of abstraction.

Sankey diagrams of the abstraction process' aggregation of treatment event sequences across the population depict the degree of concentration of treatment paths and provides different views of the data representation that could each provide valuable information to

various healthcare stakeholders. While the V2H2 level diagram can help identify patients on an adjuvant or neo-adjuvant course of chemotherapy based on the ordering of events, the more complex V1H0 level diagram shows the different paths patients take after a breast conserving surgery or mastectomy. The Sankey diagram for the re-excision rate quality measure shows the value of zooming in on a specific patient population to show the full path of surgical events versus just reporting a re-excision rate. Quality measures are often communicated by providing a rate value; visualizations such as Sankey diagrams have the potential to provide a more comprehensive view. This demonstrates the potential power of pathway exploration to provide a more holistic view of care paths for quality improvement, discovery, and educational purposes.

6.2: Informatics Limitations and Future Directions

The evaluation of our abstraction method is limited as it still relies on manually curated and structured cancer registry data. Cancer registries are often missing important information as they do not include the entire clinical population due to their inclusion criteria and do not contain all the clinical data due to the limited reporting requirements. Furthermore, the curation effort by cancer centers requires a staff of full time employees to review and code a clinical case at an estimated pace of 5 cases per day per person(Kolender, 2009). There is often a long delay (i.e., 6 months) from when care is delivered to when the case is encoded in the registry. In this encoding process, rich information in the medical record is lost. In order to achieve the goal of real-time quality improvement frameworks, we will need to use administrative data and data derived from the medical record. This will come with its own data

extraction, quality, and standardization challenges but will also provide the opportunity to represent more granular events.

Our method is also prevented from being truly scalable because we used manually derived concept mappings and an explicitly selected section of the NCI Thesaurus hierarchy. We will have to be able to not only develop our treatment mappings in a scalable way but also efficiently extract the hierarchy as needed for a particular quality question.

Our treatment event sequences maintain temporal order but the actual event times are abstracted away. This prevents queries based on time intervals or time lengths of event paths that is necessary for some certification programs' quality measures. In future work, the framework will be extended to include additional temporal relationships.

Our abstraction framework as well as our use of Sankey diagrams allows our treatment path visualizations to be scalable in handling many patients as well as handling long and variable treatment paths. Vertical abstraction collapses the types of potential treatment events, horizontal abstraction consolidates the length of the treatment paths, while the Sankey diagram aggregates sequences based on their prefix, root, and suffix. These simplified representations of complex treatment paths have the potential to provide additional visual context for treatment event sequences and can be beneficial in understanding patient populations in comparison to a single quality measure rate. The integration of an interactive visualization system into a temporal abstraction framework focused on continuous quality measurement could help improve the quality of care and shared decision-making.

6.3: Clinical Contributions

This abstraction process could be used to more systematically generate insights into clinical care. We could identify patterns in the abstracted treatment event sequences either visually with the Sankey diagram or through analytical methods. For example, one of the most common patterns at the V1Hx level is a breast conserving surgery (BCS) event followed by radiation therapy. This is, in fact, an important quality measure as the use of radiation following a BCS can reduce the rate of local cancer recurrence. Another common pattern at the V1H0 level is a BCS followed by another BCS, or a re-excision event. A number of common patterns derived from the event sequences represent important processes of care and could be used as input patterns in our framework. Additionally, the abstracted treatment sequences can be compared according to a variety of patient factors such as cancer stage, comorbidities, or whether they were referred to Vanderbilt after an initial treatment failed. For example, Figure 5.2.1 demonstrates how more frequently occurring treatment paths have a lower average cancer stage and paths involving neoadjuvant chemotherapy (chemotherapy before curative therapy) have a higher cancer stage. This can provide new ways of organizing the abstracted treatment event sequences and potentially help a user interpret infrequently occurring sequences in comparison to a similar frequent one.

Our pathway exploration tool extends this opportunity by enabling users to dynamically look at how various patient and tumor features interact with treatment paths across the VUMC or NCDB breast cancer patient population. This tool enables users to develop a more comprehensive feel for the relative frequencies of events (i.e., rate of hormone therapy in triple negative breast cancer patients) as well as avoid misinterpreting potential irregularities in the

data (i.e., HER2 data only exists after 2009). As a result, this tool can be used to answer clinical questions or educate new breast cancer patients. By incorporating the abstraction methodology into the treatment path, the user can view the treatment path at the highest level of abstraction by default. This avoids overloading the user with unnecessary data. Instead, the user can dig deeper through the vertical and horizontal abstraction controls to gain a deeper view of the treatment paths.

Our abstraction methodology can generate quality measures in a more scalable way and with less manual effort compared to existing methods. Manual abstraction from the medical record is required for QOPI and NQMBC. Prior studies have shown that the, at most, semiannual QOPI abstraction process requires ~48 to 90 minutes per patient, although the abstractor becomes more efficient with more experience(Blayney et al., 2009). This is similar to the effort required at Vanderbilt where 40 man-hours for 80 patient abstractions are required each year with an expected 100 man-hours every third year for re-certification. The RQRS is an automated system for generating specific quality measures from cancer registry data. However, RQRS only generates a few, specific measures from structured cancer registry elements. Our system can automatically generate quality measures in a more dynamic and scalable fashion. We are able to use different patterns to generate various quality measures and have the potential to take in diagnostic, clinical encounter, and outcome event data from a variety of sources (i.e., cancer registry, administrative data, medical record) in addition to treatment event data. We were able to demonstrate that our framework is generalizable enough to represent the majority of quality measures from the QOPI, NQMBC, and RQRS with the only additional requirement being the inclusion of temporal constraints.

We demonstrated how the pathway exploration platform could help users replicate

previously conducted studies by assessing the rate of use of post-mastectomy radiation therapy. Additionally, we were able to show how our system could help track the results going forward without having to repeat the study as the trend could be viewed using a future dataset. There are two primary limiting factors which need to be addressed in order to be able to handle the full breadth of clinical questions that researchers have asked of the NCDB dataset. First, the tool needs to make all of the data elements in the NCDB available through the pathway exploration tool. Many of the studies use data elements that are not currently available (i.e., specific TNM staging, insurance status, type of radiation therapy). Second, the ability to filter on the duration between treatment events will be needed to answer various temporally-related questions.

Our data-driven decision aid for new breast cancer patients demonstrates how local VUMC cancer registry data can be used to inform VUMC patients on the performance and risks of treatment at VUMC. Vanderbilt patients can understand how their specific providers tend to treat patients and develop a local assessment of care rather than national quality measures. Additionally, our decision aid tool uses a simple, natural scrolling motion to introduce potentially complex Sankey diagrams piece-by-piece in the form of a story. By allowing the patient to scroll forward and backward while they read the text and observe Sankey diagram components fading in and out, they are able to work at their own pace to develop an understanding of care. Additionally, this allows us to manage what external sites we provide links for. Patients often will search the Internet for information on their disease and potential treatment options. Providing links in this site allows us to curate the best sources of information for our patients.

6.4: Clinical Limitations and Future Directions

The visualizations generated from our abstractions have not yet been formally evaluated to assess their usefulness to various stakeholders in healthcare. Treatment path visualizations could be beneficial in a variety of contexts including quality measurement and clinical research, healthcare provider feedback, and patient education. Qualitative, observational analysis will be required to understand the needs of these stakeholders and their current workflows. Additionally, it will be important to design and build web-based applications based on our abstraction framework in an iterative fashion while receiving feedback from end users in the context of the intended use-case(Isenberg, Zuk, Collins, & Carpendale, 2008; Plaisant, 2004).

The pathway exploration platform could be delivered to cancer registries around the country. State cancer registries as well as hospitals report their data according to the North American Association for Central Cancer Registries (NAACCR) standard which is also used by the NCDB. As a result of this standard, a pathway exploration platform could be delivered as a service to various reporting institutions with little additional overhead. This service would give cancer centers more transparency into the care they provide on an ongoing basis (as the institutional cancer registry teams typically operate 6 months behind). Additionally, we could provide institution-specific data-driven decision aids for their patient population.

We plan to evaluate our data-driven decision as part of a pilot program at the Vanderbilt breast center. The decision aid will be provided to new breast cancer patients prior to their first visit with their breast surgeon. Patients will have the opportunity to review the material in order to be better informed before their meeting. We will uses a series of surveys to evaluate

the patient's opinion prior to viewing the tool, after viewing the tool, after the decision is made, and after the treatment is complete. Patients will be asked about their surgical preferences, their decision making process (make decision by self, let doctor make decision, or shared decision making), the usability of the tool, and satisfaction with the decision and treatment. We plan to extend our evaluation to assess the usability of the pathway exploration tool. This tool could allow the surgeon to personalize the filters to match their patient and show what the potential treatment options are in their case.

Users can also use a comparison tool where the user is shown two independently filterable patient populations and Sankey diagrams on the right and left of the screen. The user can then filter to select two separate patient populations and care paths. The number of patients-per-year in each of the two patient populations can then be compared either directly or as a numerator and denominator. A direct comparison can help users assess two separate populations that otherwise cannot be assessed in the same Sankey diagram. For example, this side-by-side view can help compare populations with different patient or hospital characteristics as well as significantly different care paths. Two patient populations can also be represented as a numerator and denominator to calculate the numerator populations' relative assess a specific quality measure.

REFERENCES

Abernethy, a. P., Etheredge, L. M., Ganz, P. a., Wallace, P., German, R. R., Neti, C., … Murphy, S. B. (2010). Rapid-Learning System for Cancer Care. *Journal of Clinical Oncology*, *28*(27), 4268–4274. http://doi.org/10.1200/JCO.2010.28.5478

Adlassnig, K.-P., Combi, C., Das, A. K., Keravnou, E. T., & Pozzi, G. (2006). Temporal representation and reasoning in medicine: Research directions and challenges. *Artificial Intelligence in Medicine*, *38*(2), 101–13. http://doi.org/10.1016/j.artmed.2006.10.001

Advancing Meaningful Use: Simplifying Complex Clinical Metrics Through Visual Representation. (2010). *PIIM Research*, 1–8.

Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, (May), 1–10.

Agrawal, R., & Srikant, R. (1995). Mining sequential patterns. *Data Engineering, 1995. Proceedings of …*, 3–14. http://doi.org/10.1109/ICDE.1995.380415

Allen, J. (1984). Towards a general theory of action and time. *Artificial Intelligence*, *23*.

Ancker, J. S., Kern, L. M., Edwards, A., Nosal, S., Stein, D. M., Hauser, D., & Kaushal, R. (2014). How is the electronic health record being used? Use of EHR data to assess physician-level variability in technology use. *Journal of the American Medical Informatics Association*, amiajnl-2013-002627-. http://doi.org/10.1136/amiajnl-2013-002627

Atreya, R. (2015). *Temporal Abstraction for Generating Quality Metrics for Breast Cancer Treatment*.

Augusto, J. C. (2005). Temporal reasoning for decision support in medicine. *Artificial Intelligence in Medicine*, *33*(1), 1–24. http://doi.org/10.1016/j.artmed.2004.07.006

Balaban, M., Boaz, D., & Shahar, Y. (2003). Applying Temporal Abstraction in Medical Information Systems. *Ann Math Comput Teleinform*, *1*(1), 56–64.

Bellazzi, R., Ferrazzi, F., & Sacchi, L. (2011). Predictive data mining in clinical medicine: a focus on selected methods and applications. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *1*(5), 416–430. http://doi.org/10.1002/widm.23

Bellazzi, R., Sacchi, L., & Concaro, S. (2009). Methods and tools for mining multivariate temporal data in clinical and biomedical applications. *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, *2009*, 5629–32. http://doi.org/10.1109/IEMBS.2009.5333788

Bilimoria, K. Y., Stewart, A. K., Winchester, D. P., & Ko, C. Y. (2008). The National Cancer Data Base: a powerful initiative to improve cancer care in the United States. *Annals of Surgical Oncology*, *15*(3), 683–90. http://doi.org/10.1245/s10434-007-9747-3

Black, D. M., Jiang, J., Kuerer, H. M., Buchholz, T. A., & Smith, B. D. (2014). Racial disparities in adoption of axillary sentinel lymph node biopsy and lymphedema risk in women with breast cancer. *JAMA Surgery*, *149*(8), 788–796. http://doi.org/10.1001/jamasurg.2014.23

Blayney, D. W., McNiff, K., Hanauer, D., Miela, G., Markstrom, D., & Neuss, M. (2009).

Implementation of the Quality Oncology Practice Initiative at a university comprehensive cancer center. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *27*(23), 3802–7. http://doi.org/10.1200/JCO.2008.21.6770

Bleiche, R., Ruth, K., Sigurdson, E., Beck, R., Ross, E., Wong, Y., … Al, E. (2016). Time to surgery and breast cancer survival in the united states. *JAMA Oncology*, *2*(3), 330–339. http://doi.org/10.1001/jamaoncol.2015.4508

Boaz, D., & Shahar, Y. (2005). A framework for distributed mediation of temporal-abstraction queries to clinical databases. *Artificial Intelligence in Medicine*, *34*(1), 3–24. http://doi.org/10.1016/j.artmed.2004.07.009

Boyd, C. M., Darer, J., Boult, C., Fried, L. P., Boult, L., & Wu, A. W. (2005). Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases: implications for pay for performance. *JAMA : The Journal of the American Medical Association*, *294*(6), 716–24. http://doi.org/10.1001/jama.294.6.716

Burgin, C. (2010). NAPBC accreditation gains national attention. *Journal of Registry Management*, *37*(3), 121.

Campbell, J. I., Kanters, S., Bennett, J. E., Thorlund, K., Tsai, A. C., Mills, E. J., & Siedner, M. J. (2015). Comparative Effectiveness of Induction Therapy for Human Immunodeficiency Virus-Associated Cryptococcal Meningitis: A Network Meta-Analysis. *Ofid*, *1*(Suppl 1), 1–8. http://doi.org/10.1093/o

Chakravarty, S., & Shahar, Y. (2001). Acquisition and analysis of repeating patterns in time-oriented clinical data. *Methods of Information in Medicine*, *40*(5), 410–20.

Chakravarty, S., & Shahar, Y. (2001). CAPSUL: A constraint-based specification of repeating patterns in time-oriented data. *Annals of Mathematics and Artificial Intelligence*, *30*(1–4), 3–22. http://doi.org/10.1023/A:1016661915959

Chen, A. Y., Halpern, M. T., Schrag, N. M., Stewart, A., Leitch, M., & Ward, E. (2008). Disparities and trends in sentinel lymph node biopsy among early-stage breast cancer patients (1998-2005). *Journal of the National Cancer Institute*, *100*(7), 462–474. http://doi.org/10.1093/jnci/djn057

Chen, R. C. (2013). Guideline-adherent care vs quality care in cancer patients: twins or distant cousins? *JAMA Internal Medicine*, *173*(7), 569–70. http://doi.org/10.1001/jamainternmed.2013.209

College of American Pathology (CAP). (n.d.).

College of American Radiology. (n.d.).

Collins, F. (2011). Reengineering Translational Science: The Time is Right. *Science Translational Medicine*.

Combi, C., & Shahar, Y. (1997). Temporal reasoning and temporal data maintenance in medicine: issues and challenges. *Computers in Biology and Medicine*, *27*(5), 353–68.

Commission on Cancer. (n.d.).

Conover, C. (2012). Health Care Regulation: A $169 Billion Hidden Tax. *Cato Instittute*, (527).

Czechura, T., Winchester, D. J., Pesce, C., Huo, D., Winchester, D. P., & Yao, K. (2013). Accelerated Partial-Breast Irradiation Versus Whole-Breast Irradiation for Early-Stage Breast Cancer Patients Undergoing Breast Conservation , 2003 – 2010 : A Report from the

National Cancer Data Base, 3223–3232. http://doi.org/10.1245/s10434-013-3154-8

Danciu, I., Cowan, J. D., Basford, M., Wang, X., Saip, A., Osgood, S., … Harris, P. a. (2014). Secondary use of clinical data: The Vanderbilt approach. *Journal of Biomedical Informatics*. http://doi.org/10.1016/j.jbi.2014.02.003

de Coronado, S., Haber, M. W., Sioutos, N., Tuttle, M. S., & Wright, L. W. (2004). NCI Thesaurus: using science-based terminology to integrate cancer research results. *Studies in Health Technology and Informatics*, *107*(Pt 1), 33–37.

de Zegher-Geets, I. M., Freeman, A. G., Walker, M. G., Blum, R. L., & Wiederhold, G. (1988). Summarization and display of on-line medical records. *M.D. Computing : Computers in Medical Practice*, *5*(3), 38–45, 54.

Denny, J. C. (2012). Chapter 13: Mining electronic health records in the genomics era. *PLoS Computational Biology*, *8*(12), e1002823. http://doi.org/10.1371/journal.pcbi.1002823

Desantis, C., Jemal, A., & Ward, E. (2010). Disparities in breast cancer prognostic factors by race, insurance status, and education. *Cancer Causes and Control*, *21*(9), 1445–1450. http://doi.org/10.1007/s10552-010-9572-z

Downs, S. M., Walker, M. G., & Blum, R. L. (1986). Automated summarization of on-line medical records. *Computer Science Department. Knowledge Systems Laboratory*.

Edge, S. B. (2013). The challenge of quality in breast care: beyond accreditation. *Journal of Oncology Practice / American Society of Clinical Oncology*, *9*(2), e71-2. http://doi.org/10.1200/JOP.2012.000792

Etheredge, L. M. (2007). A rapid-learning health system. *Health Affairs (Project Hope)*, *26*(2), w107-18. http://doi.org/10.1377/hlthaff.26.2.w107

Etheredge, L. M. (2014). Rapid learning: a breakthrough agenda. *Health Affairs (Project Hope)*, *33*(7), 1155–62. http://doi.org/10.1377/hlthaff.2014.0043

Fagan, L., & Kunz, J. (1984). Extensions to the Rule-Based Formalism for a Monitoring Task. *Rule-Based Expert …*, 397–423.

Fedewa, S. A., Ward, E. M., Stewart, A. K., & Edge, S. B. (2010). Delays in adjuvant chemotherapy treatment among patients with breast cancer are more likely in African American and Hispanic populations: A national cohort study 2004-2006. *Journal of Clinical Oncology*, *28*(27), 4135–4141. http://doi.org/10.1200/JCO.2009.27.2427

Fedewa, S. a, Edge, S. B., Stewart, A. K., Halpern, M. T., Marlow, N. M., & Ward, E. M. (2011). Race and ethnicity are associated with delays in breast cancer treatment (2003-2006). *Journal of Health Care for the Poor and Underserved*, *22*(1), 128–141. http://doi.org/10.1353/hpu.2011.0006

Freedman, R. A., Virgo, K. S., He, Y., Pavluck, A. L., Winer, E. P., Ward, E. M., & Keating, N. L. (2011). The association of race/ethnicity, insurance status, and socioeconomic factors with breast cancer care. *Cancer*, *117*(1), 180–189. http://doi.org/10.1002/cncr.25542

Freedman, R. A., Virgo, K. S., Labadie, J., He, Y., Partridge, A. H., & Keating, N. L. (2012). Receipt of locoregional therapy among young women with breast cancer. *Breast Cancer Research and Treatment*, *135*(3), 893–906. http://doi.org/10.1007/s10549-012-2156-8

Friedman, C. P. (2009). A “fundamental theorem” of biomedical informatics. *Journal of the American Medical Informatics Association : JAMIA*, *16*(2), 169–70.

http://doi.org/10.1197/jamia.M3092

Gesme, D. H., & Wiseman, M. (2011). Strategic Use of Clinical Pathways. *Journal of Oncology Practice*, *7*(1), 54–56. http://doi.org/10.1200/JOP.2010.000193

Goldwein, J. W., & Rose, C. M. (2007). QOPI, EHRs, and Quality Measurement. *Journal of Oncology Practice / American Society of Clinical Oncology*, *3*(6), 340. http://doi.org/10.1200/JOP.0769001

Greenfield, D., & Braithwaite, J. (2008). Health sector accreditation research: a systematic review. *International Journal for Quality in Health Care : Journal of the International Society for Quality in Health Care / ISQua*, *20*(3), 172–83. http://doi.org/10.1093/intqhc/mzn005

Greif, J. M., Pezzi, C. M., Klimberg, V. S., Bailey, L., & Zuraek, M. (2012). Gender differences in breast cancer: analysis of 13,000 breast cancers in men from the National Cancer Data Base. *Annals of Surgical Oncology*, *19*(10), 3199–204. http://doi.org/10.1245/s10434-012-2479-z

Guerra-Gomez, J., Pack, M., Plaisant, C., & Shneiderman, B. (2013). Visualizing Change over Time Using Dynamic Hierarchies: TreeVersity2 and the StemView. *Visualization and Computer Graphics, IEEE Transactions on*, *19*(12), 2566–2575.

Haimowitz, I. J., & Kohane, I. S. (1993). An epistemology for clinically significant trends, 176–181.

Haimowitz, I., & Kohane, I. (1993). Automated trend detection with alternate temporal hypotheses. *IJCAI*, 146–151.

Halpern, M., Chen, A., Marlow, N., & Ward, E. (2009). Disparities in Receipt of Lymph Node Biopsy Among Early-Stage Female Breast Cancer Patients. *Annals of Surgical Oncology*, *16*(3), 562–570. http://doi.org/10.1245/s10434-008-0205-7

Halpern, M. T., Bian, J., Ward, E. M., Schrag, N. M., & Chen, A. Y. (2007). Insurance status and stage of cancer at diagnosis among women with breast cancer. *Cancer*, *110*(2), 403–411. http://doi.org/10.1002/cncr.22786

Hamburg, M. A., & Collins, F. S. (2010). The path to personalized medicine. *The New England Journal of Medicine*, *363*(4), 301–4.

Hou, N., Yao, K., Jaskowiak, N. T., Hasan, Y., & Winchester, D. J. (2012). Patterns of regional node irradiation therapy following breast - conserving surgery : Report from the National Cancer Data, *30*(15), 2016.

Hoverman, J. R., Cartwright, T. H., Patt, D. A., Espirito, J. L., Clayton, M. P., Garey, J. S., … Beveridge, R. A. (2011). Pathways, outcomes, and costs in colon cancer: Retrospective evaluations in two distinct databases. *Journal of Oncology Practice*, *7*(3 SUPPL.), 52s–59s. http://doi.org/10.1200/JOP.2011.000318

Huo, D., Hou, N., Jaskowiak, N. T., Hasan, Y., Winchester, D. J., Winchester, D. P., & Yao, K. (2012). Postmastectomy radiotherapy for patients with one to three positive lymph nodes: Utilization and benefit. *Journal of Clinical Oncology*, *30*(15), 2012.

IBM Netezza Data Warehouse Appliances – The Simple Data Warehouse Appliance for Serious Analytics. (2014, April 1). IBM Corporation.

In, H., Bilimoria, K. Y., Stewart, A. K., Wroblewski, K. E., Posner, M. C., Talamonti, M. S., & Winchester, D. P. (2014). Cancer Recurrence: An Important but Missing Variable in National Cancer Registries. *Annals of Surgical Oncology*, *21*(5), 1520–1529.

http://doi.org/10.1245/s10434-014-3516-x

Isenberg, P., Zuk, T., Collins, C., & Carpendale, S. (2008). Grounded Evaluation of Information Visualization. *Proceedings of the 2008 Conference on BEyond Time and Errors Novel evaLuation Methods for Information Visualization - BELIV '08*, 1. http://doi.org/10.1145/1377966.1377974

James, B. C., & Savitz, L. A. (2011). How Intermountain trimmed health care costs through robust quality improvement efforts. *Health Affairs (Project Hope)*, *30*(6), 1185–91. http://doi.org/10.1377/hlthaff.2011.0358

K., Y., T., C., D., H., D.P., W., S.B., E., & D.J., W. (2012). Utilization of post-lumpectomy radiation therapy in women 70 years of age or older: A report from the National Cancer Data Base. *Journal of Clinical Oncology*, *30*(15 SUPPL. 1), 9343.

Kahn, M. G. (1991). Modeling time in medical decision-support programs. *Medical Decision Making : An International Journal of the Society for Medical Decision Making*, *11*(4), 249–64.

Kahn, M. G., Fagan, L. M., & Sheiner, L. B. (1991). Combining physiologic models and symbolic methods to interpret time-varying patient data. *Methods of Information in Medicine*, *30*(3), 167–78.

Kahn, M. G., Fagan, L. M., & Tu, S. (1991). Extensions to the time-oriented database model to support temporal reasoning in medical expert systems. *Methods of Information in Medicine*, *30*(1), 4–14.

Kahn, M., Tu, S., & Fagan, L. (1991). TQuery: A Context-Sensitive temporal query language. *Computers and Biomedical Research*.

Kaufman, C. S., Shockney, L., Rabinowitz, B., Coleman, C., Beard, C., Landercasper, J., … Wiggins, D. (2010). National Quality Measures for Breast Centers (NQMBC): a robust quality tool: breast center quality measures. *Annals of Surgical Oncology*, *17*(2), 377–385. http://doi.org/10.1245/s10434-009-0729-5

Keller, D., Guilfoyle, C., & Sariego, J. (2011). Geographical influence on racial disparity in breast cancer presentation in the United States. *American Surgeon*, *77*(7), 933–936.

Kennedy, T., Stewart, A. K., Bilimoria, K. Y., Patel-Parekh, L., Sener, S. F., & Winchester, D. P. (2007). Treatment trends and factors associated with survival in T1aN0 and T1bN0 breast cancer patients. *Annals of Surgical Oncology*, *14*(10), 2918–2927. http://doi.org/10.1245/s10434-007-9441-5

Killelea, B. K., Yang, V. Q., Wang, S.-Y., Hayse, B., Mougalian, S., Horowitz, N. R., … Lannin, D. R. (2015). Racial Differences in the Use and Outcome of Neoadjuvant Chemotherapy for Breast Cancer: Results From the National Cancer Data Base. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *33*(36), 4267–76. http://doi.org/10.1200/JCO.2015.63.7801

Kinsman, L., Rotter, T., James, E., Snow, P., & Willis, J. (2010). What is a clinical pathway? Development of a definition to inform the debate. *BMC Medicine*, *8*, 31. http://doi.org/10.1186/1741-7015-8-31

Klimov, D., Shahar, Y., & Taieb-Maimon, M. (2010). Intelligent visualization and exploration of time-oriented data of multiple patients. *Artificial Intelligence in Medicine*, *49*(1), 11–31. http://doi.org/10.1016/j.artmed.2010.02.001

Knops, A., Legemate, D., Md, P., Goossens, A., Bossuyt, P., & Ubbink, D. (2013). Decision Aids for Patients Facing a Surgical Treatment Decision: A Systematic Review and Meta-analysis. *Ann Surg*, *257*(5), 860–866. http://doi.org/10.1097/SLA.0b013e3182864fd6

Kohane, I., & Haimowitz, I. (1993). Hypothesis-driven data abstraction with trend templates. *… of the Annual Symposium on Computer …*.

Kolender, E. (2009). Picture perfect staffing in today's economy. *Journal of Registry Management*, 153–154.

Kowalski, R., & Sergot, M. (1989). A logic-based calculus of events. *Foundations of Knowledge Base Management*, *4*(June 1975).

Kummerow, K. L., Du, L., Penson, D. F., Shyr, Y., & Hooks, M. a. (2015). Nationwide trends in mastectomy for early-stage breast cancer. *JAMA Surgery*, *150*(1), 9–16. http://doi.org/10.1001/jamasurg.2014.2895

Mariotto, A. B., Yabroff, K. R., Shao, Y., Feuer, E. J., & Brown, M. L. (2011). Projections of the cost of cancer care in the United States: 2010-2020. *Journal of the National Cancer Institute*, *103*(2), 117–28. http://doi.org/10.1093/jnci/djq495

Mcdermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science*, *6*(2), 101–155. http://doi.org/10.1016/S0364-0213(82)90003-9

Merkow, R. P., Chung, J. W., Paruch, J. L., Bentrem, D. J., & Bilimoria, K. Y. (2014). Relationship between cancer center accreditation and performance on publicly reported quality measures. *Annals of Surgery*, *259*(6), 1091–7. http://doi.org/10.1097/SLA.0000000000000542

Miksch, S., Shahar, Y., & Johnson, P. (1997). ASBRU : A TASK-SPECIFIC , INTENTION-BASED , AND TIME-ORIENTED LANGUAGE FOR REPRESENTING SKELETAL PLANS. *… of the 7th Workshop on Knowledge …*, 1–25.

Miller, R. A., & Masarie, F. E. (1990). The demise of the "Greek Oracle" model for medical diagnostic systems. *Methods of Information in Medicine*, *29*(1), 1–2.

Minami, C. A., Bilimoria, K. Y., Hansen, N. M., Strauss, J. B., Hayes, J. P., Feinglass, J. M., … Yang, A. D. (2016). National Evaluation of the New Commission on Cancer Quality Measure for Postmastectomy Radiation Treatment for Breast Cancer. *Annals of Surgical Oncology*. http://doi.org/10.1245/s10434-016-5257-5

Monroe, M., Lan, R., Lee, H., Plaisant, C., & Shneiderman, B. (2013). Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics*, *19*(12), 2227–2236. http://doi.org/10.1109/TVCG.2013.200

Moskovitch, R., & Shahar, Y. (2009). Medical Temporal-Knowledge Discovery via Temporal Abstraction. *AMIA Annual Symposium Proceedings*, 452–456.

Musen, M. A., Tu, S. W., Das, A. K., & Shahar, Y. (1996). EON: A Component-Based Approach to Automation of Protocol-Directed Therapy. *Journal of the American Medical Informatics Association*, *3*(6), 367–388.

NAACCR Chapter III: Standards for Tumor Inclusion and Reportability. (n.d.). Retrieved November 3, 2015, from http://www.naaccr.org/Applications/ContentReader/Default.aspx?c=3

NAACCR Chapter X: Data Dictionary. (n.d.). Retrieved November 3, 2015, from

http://www.naaccr.org/Applications/ContentReader/Default.aspx?c=10

National Accreditation Program for Breast Centers (NAPBC). (n.d.).

National Quality Forum (NQF). (n.d.).

National Quality Measures for Breast Centers (NQMBC). (n.d.).

NCI Thesaurus. (n.d.). Retrieved from http://ncit.nci.nih.gov/

Neubauer, M. A., Hoverman, J. R., Kolodziej, M., Reisman, L., Gruschkus, S. K., Hoang, S., … Beveridge, R. A. (2010). Cost Effectiveness of Evidence-Based Treatment Guidelines for the Treatment of Non-Small-Cell Lung Cancer in the Community Setting. *Journal of Oncology Practice*, *6*(1), 12–18. http://doi.org/10.1200/JOP.091058

Newhouse, J. P., Garber, A. M., Graham, R. P., Margaret, A., Mancher, M., & Kibria, A. (2013). *Variation in Health Care Spending : Target Decision Making , Not Geography*.

Newman, L. A., Lee, C. T., Parekh, L. P., Stewart, A. K., Thomas, C. R., Beltran, R. A., … Nelson, H. (2006). Use of the National Cancer Data Base to develop clinical trials accrual targets that are appropriate for minority ethnicity patients: A report from the American College of Surgeons Oncology Group (ACOSOG) special populations committee. *Cancer*, *106*(1), 188–195. http://doi.org/10.1002/cncr.21592

Nguyen, J., Shahar, Y., & Tu, S. (1999). Integration of temporal reasoning and temporal-data maintenance into a reusable database mediator to answer abstract, time-oriented queries: The Tzolkin system. *… Information Systems*, *145*, 121–145.

Norén, G. N., Hopstadius, J., Bate, A., Star, K., & Edwards, I. R. (2009). Temporal pattern discovery in longitudinal electronic patient records. *Data Mining and Knowledge Discovery*, *20*(3), 361–387. http://doi.org/10.1007/s10618-009-0152-3

Orphanou, K., Stassopoulou, A., & Keravnou, E. (2014). Temporal abstraction and temporal Bayesian networks in clinical domains: a survey. *Artificial Intelligence in Medicine*, *60*(3), 133–49. http://doi.org/10.1016/j.artmed.2013.12.007

Oshima Lee, E., & Emanuel, E. J. (2013). Shared decision making to improve care and reduce costs. *The New England Journal of Medicine*, *368*(1), 6–8. http://doi.org/10.1056/NEJMp1209500

Peterson, J. (2012). Implementing Quality Oncology Practice Initiative (QOPI) participation in a community oncology practice. *Journal of Registry Management*, *39*(4), 154–7.

Pezzi, C. M., Patel-Parekh, L., Cole, K., Franko, J., Klimberg, V. S., & Bland, K. (2007). Characteristics and treatment of metaplastic breast cancer: analysis of 892 cases from the National Cancer Data Base. *Annals of Surgical Oncology*, *14*(1), 166–173. http://doi.org/10.1245/s10434-006-9124-7

Pierce, O., & Allen, M. (2015). Assessing surgeon-level risk of patient harm during elective surgery for public reporting. *ProPublica*, 1–21.

Plaisant, C. (2004). The Challenge of Information Visualization Evaluation.

Plaisant, C., Milash, B., Rose, A., Widoff, S., & Shneiderman, B. (1996). LifeLines: visualizing personal histories. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Common Ground - CHI '96*, 221-. http://doi.org/10.1145/238386.238493

Poonacha, T. K., & Go, R. S. (2011). Level of scientific evidence underlying recommendations

arising from the National Comprehensive Cancer Network clinical practice guidelines. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *29*(2), 186–91. http://doi.org/10.1200/JCO.2010.31.6414

Porter, M. E. (2010). What is value in health care? *The New England Journal of Medicine*, *363*(26), 2477–81. http://doi.org/10.1056/NEJMp1011024

Powsner, S., & Tufte, E. (1994). Graphical summary of patient status. *The Lancet*.

Quality Oncology Practice Initiative (QOPI). (n.d.).

Quill, T. E., & Holloway, R. G. (2012). Evidence, preferences, recommendations--finding the right balance in patient care. *The New England Journal of Medicine*, *366*(18), 1653–5. http://doi.org/10.1056/NEJMp1201535

RA, M., HE, P. J., & JM, M. (1982). INTERNIST-1, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. *New Eng J Med*.

Raj, R., O'Connor, M., & Das, A. (2007). An ontology-driven method for hierarchical mining of temporal patterns: application to HIV drug resistance research. *AMIA Annual Symposium …*, 614–619.

Ravdin, P. (2001). Computer program to assist in making decisions about adjuvant therapy for women with early breast cancer. *J Clin Oncol*, *19*(4), 980–991.

Riehmann, P., Hanfler, M., & Froehlich, B. (2005). Interactive sankey diagrams. *Proceedings - IEEE Symposium on Information Visualization, INFO VIS*, 233–240. http://doi.org/10.1109/INFVIS.2005.1532152

Rueth, N. M., Lin, H. Y., Bedrosian, I., Shaitelman, S. F., Ueno, N. T., Shen, Y., & Babiera, G. (2014). Underuse of trimodality treatment affects survival for patients with inflammatory breast cancer: An analysis of treatment and survival trends from the National Cancer Database. *Journal of Clinical Oncology*, *32*(19), 2018–2024. http://doi.org/10.1200/JCO.2014.55.1978

Russ, T. a. (1995). Use of data abstraction methods to simplify monitoring. *Artificial Intelligence in Medicine*, *7*(6), 497–514. http://doi.org/10.1016/0933-3657(95)00024-8

Shahar, Y. (2013). The Elicitation, Representation, Application, and Automated Discovery of Time-Oriented Declarative Clinical Knowledge. *Process Support and Knowledge Representation in …*, 1–29.

Shahar, Y., Goren-Bar, D., Boaz, D., & Tahan, G. (2006). Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artificial Intelligence in Medicine*, *38*(2), 115–35. http://doi.org/10.1016/j.artmed.2005.03.001

Shahar, Y., Miksch, S., & Johnson, P. (1998). The Asgaard Project : A Task-Specific Framework for the Application and Critiquing of Time-Oriented Clinical Guidelines. *Artificial Intelligence in Medicine*, 1–25.

Shahar, Y., & Musen, M. (1992). RESUME: A Temporal-Abstraction System for Patient Monitoring. *Proceedings of the Annual Symposium on …*.

Shahar, Y., & Musen, M. A. (1996). Knowledge-based temporal abstraction in clinical domains. *Artificial Intelligence in Medicine*, *8*(3), 267–298.

Shahar, Y., Tu, S., & Musen, M. (1992). Knowledge Acquisition for Temporal-Abstraction Mechanisms. *Knowledge Acquisition*.

Sharpe, S. M., Liederbach, E., Czechura, T., Pesce, C., Winchester, D. J., & Yao, K. (2014). Impact of bilateral versus unilateral mastectomy on short term outcomes and adjuvant therapy, 2003-2010: a report from the national cancer data base. *Annals of Surgical Oncology*, *21*(9), 2920–7. http://doi.org/10.1245/s10434-014-3687-5

Shoham, Y. (1987). Temporal logics in AI: Semantical and ontological considerations. *Artificial Intelligence*, *33*(1), 89–104. http://doi.org/10.1016/0004-3702(87)90052-X

Shortliffe, E. H. (1977, October 5). Mycin: A Knowledge-Based Computer Program Applied to Infectious Diseases. *Proceedings of the Annual Symposium on Computer Application in Medical Care*. American Medical Informatics Association.

Siegel, R., Ma, J., Zou, Z., & Jemal, A. (2014). Cancer Statistics, 2014. *CA: A Cancer Journal for Clinicians*, *64*(1), 9–29. http://doi.org/10.3322/caac.21208.

Sisco, M., Du, H., Warner, J. P., Howard, M. a., Winchester, D. P., & Yao, K. (2012). Have we expanded the equitable delivery of postmastectomy breast reconstruction in the new millennium? Evidence from the national Cancer data base. *Journal of the American College of Surgeons*, *215*(5), 658–667. http://doi.org/10.1016/j.jamcollsurg.2012.07.008

Sledge, G. W., Hudis, C. a, Swain, S. M., Yu, P. M., Mann, J. T., Hauser, R. S., & Lichter, A. S. (2013). ASCO's approach to a learning health care system in oncology. *Journal of Oncology Practice / American Society of Clinical Oncology*, *9*(3), 145–8. http://doi.org/10.1200/JOP.2013.000957

Sledge, G. W., Miller, R. S., & Hauser, R. (2013). CancerLinQ and the future of cancer care. *American Society of Clinical Oncology Educational Book / ASCO. American Society of Clinical Oncology. Meeting*, 430–434. http://doi.org/10.1200/EdBook_AM.2013.33.430

Spence, I. (2006). William Playfair and the psychology of graphs. *2006 JSM Proceedings, American Statistical …*, 2426–2436.

Spokoiny, A., & Shahar, Y. (2003). Momentum – an Active Time-Oriented Database for Intelligent Abstraction , Exploration and Analysis of Clinical Data. *… Intelligent Data Analysis in Medicine and …*.

Stacey, M., & McGregor, C. (2007). Temporal abstraction in intelligent clinical data analysis: a survey. *Artificial Intelligence in Medicine*, *39*(1), 1–24. http://doi.org/10.1016/j.artmed.2006.08.002

Swan, J., Wingo, P., Clive, R., West, D., Miller, D., Hutchison, C., … Edwards, B. K. (1998). Cancer surveillance in the U.S.: can we have a national system? *Cancer*, *83*(7), 1282–91.

Swan, M. (2012). Health 2050: The Realization of Personalized Medicine through Crowdsourcing, the Quantified Self, and the Participatory Biocitizen. *Journal of Personalized Medicine*, *2*(4), 93–118. http://doi.org/10.3390/jpm2030093

Thompson, W. K., Rasmussen, L. V, Pacheco, J. a, Peissig, P. L., Denny, J. C., Kho, a N., … Pathak, J. (2012). An Evaluation of the NQF Quality Data Model for Representing Electronic Health Record Driven Phenotyping Algorithms. *AMIA Annu Symp Proc*, *2012*, 911–920.

Tinetti, M. E., Fried, T. R., & Boyd, C. M. (2012). Designing health care for the most common chronic condition--multimorbidity. *JAMA*, *307*(23), 2493–4. http://doi.org/10.1001/jama.2012.5265

Waljee, J. F., Rogers, M. a M., & Alderman, A. K. (2007). Decision aids and breast cancer: Do they influence choice for surgery and knowledge of treatment options? *Journal of Clinical*

Oncology, *25*(9), 1067–1073. http://doi.org/10.1200/JCO.2006.08.5472

Wang, E. H., Mougalian, S. S., Soulos, P. R., Rutter, C. E., Evans, S. B., Haffty, B. G., … Yu, J. B. (2014). Adoption of hypofractionated whole-breast irradiation for early-stage breast cancer: A national cancer data base analysis. *International Journal of Radiation Oncology Biology Physics*, *90*(5), 993–1000. http://doi.org/10.1016/j.ijrobp.2014.06.038

Warren, J. L., Klabunde, C. N., Schrag, D., Bach, P. B., & Riley, G. F. (2001). Overview of the SEER-Medicare Data. *Medical Care*, *40*(8), 1–19.

Weldon, C. B., Trosman, J. R., Gradishar, W. J., Benson, A. B., & Schink, J. C. (2012). Barriers to the use of personalized medicine in breast cancer. *Journal of Oncology Practice / American Society of Clinical Oncology*, *8*(4), e24-31. http://doi.org/10.1200/JOP.2011.000448

West, V. L., Borland, D., & Hammond, W. E. (2014). Innovative information visualization of electronic health record data: a systematic review. *Journal of the American Medical Informatics Association : JAMIA*, 1–7. http://doi.org/10.1136/amiajnl-2014-002955

Williams, R. T., Yao, K., Stewart, A. K., Winchester, D. J., Turk, M., Gorchow, A., … Winchester, D. P. (2011). Needle versus excisional biopsy for noninvasive and invasive breast cancer: report from the National Cancer Data Base, 2003-2008. *Annals of Surgical Oncology*, *18*(13), 3802–10. http://doi.org/10.1245/s10434-011-1808-y

Wingo, P. A., Jamison, P. M., Hiatt, R. A., Weir, H. K., Gargiullo, P. M., Hutton, M., … Hall, H. I. (2003). Building the infrastructure for nationwide cancer surveillance and control--a comparison between the National Program of Cancer Registries (NPCR) and the Surveillance, Epidemiology, and End Results (SEER) Program (United States). *Cancer Causes & Control : CCC*, *14*(2), 175–93.

Wingo, P., & Howe, H. (2005). A national framework for cancer surveillance in the United States. *Cancer Causes & …*, 151–170.

Wiznia, L. E., Lannin, D. R., Evans, S. B., Hofstatter, E. W., Horowitz, N. R., Killelea, B. K., … Chagpar, A. B. (2014). The Number of Lymph Nodes Dissected in Breast Cancer Patients Influences the Accuracy of Prognosis. *Annals of Surgical Oncology*, *21*(2), 389–394. http://doi.org/10.1245/s10434-013-3308-8

Wongsuphasawat, K. (2011). LifeFlow: visualizing an overview of event sequences. *Proceedings of the …*, 1747–1756.

Wongsuphasawat, K., & Gotz, D. (2012). Exploring Flow, Factors, and Outcomes of Temporal Event Sequences with the Outflow Visualization. *IEEE Transactions on Visualization and Computer Graphics*, *18*(12), 2659–2668. http://doi.org/10.1109/TVCG.2012.225

Yao, K., Czechura, T., Liederbach, E., Winchester, D. J., Pesce, C., Shaikh, A., … Huo, D. (2014). Utilization of Accelerated Partial Breast Irradiation for Ductal Carcinoma In Situ , 2003 – 2011 : Report from the National Cancer Database, 3457–3465. http://doi.org/10.1245/s10434-014-3717-3

Yao, K., Stewart, A. K., Winchester, D. J., & Winchester, D. P. (2010). Trends in contralateral prophylactic mastectomy for unilateral cancer: a report from the National Cancer Data Base, 1998-2007. *Annals of Surgical Oncology*, *17*(May), 2554–2562. http://doi.org/10.1245/s10434-010-1091-3

Zaki, M. (2001). SPADE: An efficient algorithm for mining frequent sequences. *Machine Learning*, 31–60.

Zhang, Z., Wang, B., Ahmed, F., & Ramakrishnan, I. (2013). The Five W's for Information Visualization with Application to Healthcare Informatics. *IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS*, *19*(11), 1895–1910.

Zon, R. T., Frame, J. N., Neuss, M. N., Page, R. D., Wollins, D. S., Stranne, S., & Bosserman, L. D. (2016). American Society of Clinical Oncology Policy Statement on Clinical Pathways in Oncology. *Journal of Oncology Practice / American Society of Clinical Oncology*, JOP.2015.009134-. http://doi.org/10.1200/JOP.2015.009134