TEMPORAL ABSTRACTION FOR GENERATING QUALITY METRICS FOR BREAST CANCER TREATMENT

By

Ravi Viswanathan Atreya

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Biomedical Informatics

August, 2015

Nashville, Tennessee

Approved:

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

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

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

ACKNOWLEDGEMENTS

I would first like to thank my advisor, Dr. Mia Levy, for her mentorship and support through this Masters work. I have enjoyed our discussions on the role of informatics in oncology and specifically breast cancer care. I would also like to thank Dr. Tom Lasko for his advice on data mining and analysis; I have enjoyed learning about new ways to think about clinical data. I would also like to thank Dr. Josh Denny for his advice on phenotyping methods. I look forward to continuing this work for the PhD.

I would also like to thank the Department of Biomedical Informatics (DBMI) and Medical Scientist Training Program (MSTP) for their support. The DBMI has provided a great environment for my research training and I would like to thank so many of the faculty for always being willing to provide advice or discuss ideas. I would specifically like to thank Drs. Cindy Gadd and Kevin Johnson for their support through the training program. I would also like to thank the staff who make everything possible, especially Rischelle Jenkins. I would also like to thank the MSTP for providing clinical, research, and leadership training. I would like to specifically thank Dr. Terry Dermody, the entire MSTP leadership team, the advising college leadership, and especially Melissa Krasnove. I am very grateful for the funding that makes my work possible, including the Vanderbilt MSTP T32 and CTSA TL1.

Finally, I would like to thank my family and friends through this process. I appreciate the friendship of my fellow DBMI, MSTP, and medical school students. I would like to thank my family, especially my parents for their guidance and support and my sister.

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

Monitoring quality metrics continuously in breast cancer care can help healthcare providers and organizations engage with patients, advance the delivery of care, and tackle financial challenges. In this work, we develop the Pathfinder method to identify treatment event patterns that relate to quality metrics in breast cancer clinical data. We used manually curated cancer registry data and administrative data to evaluate our method and compare the data sources. We aim to demonstrate that we can effectively track quality metrics by abstracting raw data to the proper level of granularity.

We developed the Pathfinder method that consists of six subtasks: data extraction, data standardization, vertical abstraction, horizontal abstraction, quality metric querying, and quality metric generation. We used cancer registry treatment event and administrative CPT data from Vanderbilt University Medical Center from 2000-2012. We characterized the data sources, assessed the abstraction process, and measured the quality of the CPT codes. We then used the data to evaluate three quality metrics: rate of re-excision after initial breast conserving surgery, radiation therapy after breast conserving surgery, and chemotherapy usage in early stage disease. Finally, we used an event sequence mining method to identify common treatment event patterns to characterize and compare our data sources.

Cancer registry and CPT data for 2679 breast cancer patients were used in our study. The application of our variable abstraction process produced an approximately 12-fold reduction in the number of unique treatment event sequences from the raw sequence to most abstracted state. The quality metrics developed from the cancer registry data matched expected national rates. The CPT data was often aligned with the curated cancer registry

events, but did show gaps with 37% of patients missing at least one CPT code. Despite this, the CPT data still produced similar surgical quality metrics to the cancer registry data. The CPT and cancer registry data did occasionally have different rates of frequent event patterns occurring in the patient population.

This work demonstrates how our temporal abstraction method can enable us to transform raw clinical data to the level of abstraction necessary to generate the desired quality metric. We were able to measure a set of quality metrics over a 12-year period with both manually curated cancer registry data and administrative CPT data that matched expected national rates. Despite this, near real-time metrics are difficult to achieve given the manual nature of cancer registry curation and the high frequency of missing CPT data from care delivered outside the organization. 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 clinical quality metrics can help healthcare overcome its many challenges, specifically in breast cancer care. Section 2.1 reviews how continuous quality measurement 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 how increased awareness of breast cancer has led to increased efforts to monitor the quality of care. Finally,

section 2.4 describes the challenges of achieving near real time quality metrics 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 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. Finally, section 3.4 discusses the hypothesis and aim of this study to track quality metrics using components of a temporal abstraction 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 six subtasks of the Pathfinder method. Section 4.3 focuses on the extraction and representation of the clinical data. Section 4.4 describes how quality metrics are generated from cancer registry treatment event data. Section 4.5 describes how administrative data quality will be assessed and how quality metrics will be generated using administrative data.

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 resulting quality metrics developed from the cancer registry data. Section 5.4 describes the quality of the surgical CPT codes by comparing them with the cancer registry data. Finally, section 5.5 describes the quality metrics developed from CPT data and compares it to that of the cancer registry data.

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 and the related limitations and future directions. Section 6.2 the same for the clinical domain.

Chapter 2: Role of Quality Metrics in Breast Cancer Care

High levels of healthcare spending along with a lack of high quality patient care have driven interest towards the delivery of high value care. The rapidly shifting healthcare landscape along with the variability of clinical care has made it challenging to track clinical quality, especially in near-real time (Section 2.1). This need is especially apparent in the care of cancer, a high mortality and high cost disease. Registries have focused on tracking cancer cases through manual reporting and a number of cancer organizations have developed quality standards. This emphasis on cancer care has driven research and clinical trials that have led to many new diagnostics and therapeutics (Section 2.2). There have been visible efforts to increase research and improve care for breast cancer, leading to an increase in the number of accreditation organizations tracking screening, diagnosis, treatment, and surveillance via quality metrics (Section 2.3). Measuring quality metrics 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. Electronic medical record (EMR) data along with new algorithms and visualizations can enable the near-real time tracking of clinical quality not only for quality reporting but also for daily care improvement (Section 2.4).

2.1: Imperative for Tracking Clinical Quality Metrics

High levels of healthcare spending in the United States persist without demonstrating a commensurate return on investment. Health care spending in 2011 was $2.7 trillion and accounted for 17.9% of gross domestic product (GDP) as well as 15.7% of US workforce (Moses et al., 2013). Much of the increase in healthcare spending is due to Medicare (and other federal healthcare spending) that consumed 20% of federal expenditures in 2009 and is a major driver

of structural federal deficits (Chernew, Baicker, & Hsu, 2010; Fuchs, 2013; Moses et al., 2013). This level of monetary and human capital investment in healthcare is not reflected in the quality of care delivery. While awareness has improved since two landmark Institute of Medicine (IOM) reports ("Crossing the Quality Chasm: A New Health System for the 21st Century," 2001; IOM, 2000) and excellence has been shown for a few quality metrics, the US has been unable to scale significant improvement across the spectrum of healthcare delivery (Chassin & Loeb, 2011; McGlynn et al., 2003). There has been increased focus on improving the value in the healthcare sector instead of the volume of services delivered. 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). New payment models such as bundled payments and accountable care organizations (ACO) are being leveraged to shift incentive structures from services delivered to outcomes achieved (C. Chen & Ackerly, 2014; Cutler & Ghosh, 2012; Schneider, Hussey, & Schnyer, 2011; Shrank, 2013). In order to monitor our shifting healthcare landscape and understand the value of delivered care, it is vital to track important indicators of clinical quality in near real-time.

As best practices in healthcare continue to evolve and advance, 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). In addition, clinical trial research systems are being developed to allow scientists to generate knowledge at a more rapid pace by lowering barriers to data collection and dissemination (Etheredge, 2014). However, studies have demonstrated a significant lag between clinical discovery and widespread implementation in practice (L. W. Green, Ottoson, García, & Hiatt, 2009; Morris, Wooding, & Grant, 2011). Not only do we need to enhance the

diffusion of new clinical knowledge, but we must also be able to quickly determine the outcome and value of new interventions. New tools to track progress in near real-time will be necessary to achieve that goal (Etheredge, 2007, 2014).

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 quality metrics from the perspective of the individual patient case.

Near real-time tracking of the quality of care will be an essential component of efforts to improve the value of healthcare by following best-practice and removing unnecessary variability from the system. Quality metrics that measure the performance and value of care are valuable for all stakeholders, including patients, providers, administrators, scientists, entrepreneurs, and

policy makers. Continuous measurement of health care quality and efficiency provides a real opportunity to make rapid improvements and adjustments in health care delivery (Chassin, Loeb, Schmaltz, & Wachter, 2010; A. Higgins, Veselovskiy, & McKown, 2013; Shrank, 2013). As more quality metrics are developed, it will be important to guard against misrepresentation of metrics, while still enabling flexibility for different measurements in various situations that hold value for the end user (Haut & Pronovost, 2011; A. Higgins et al., 2013; Meyer et al., 2012; Shahian, Wolf, Iezzoni, Kirle, & Normand, 2010). In addition, it will be important to factor how payment models affect the billing process and understand biases in administrative data sets (Farmer, Black, & Bonow, 2013). Research on clinical quality metrics is necessary to help all stakeholders make the best possible decisions in a rapidly changing healthcare environment and drive continuous quality improvement.

2.2: Importance of Quality Metrics in Cancer Care

Cancer is a highly prevalent set of diseases in which improperly regulated cells divide and spread. In 2014, there will be an estimated 1.6 million new cancer cases and approximately 0.5 million cancer related deaths in the US (Siegel, Ma, Zou, & Jemal, 2014). Cancer is now the leading cause of death for men and women between 40 and 79 years of age (Siegel et al., 2014). As of 2012, there were approximately 13.7 million cancer survivors living in the US and this is expected to increase to 18 million by 2022 (de Moor et al., 2013). Given the high prevalence, growing number of survivors, and intensive course of diagnostics and therapeutics, cancer care is costly and comprises 5% of US health care costs and 10% of Medicare expenditures (Stockdale & Guillory, 2013; Sullivan et al., 2011). In addition, costs are expected to increase by 27% from $124 billion in 2010 to $155 billion in 2020 (Mariotto, Yabroff, Shao, Feuer, & Brown, 2011). Medicare beneficiaries make up 54% of all new cancer cases (Stockdale & Guillory, 2013), and

with the aging US population, cancer costs will continue to be an important part of healthcare and fiscal policy.

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 at the National Cancer Institute (SEER at NCI). SEER collects data on cancer cases from select states and metropolitan areas. To expand the collection of data, the National Program of Cancer Registries (NPCR) overseen by the Center for Disease Control (CDC) was established through the Cancer Registries Amendment Act of 1992 to expand statewide cancer registry programs. While the development of the registries has led to significant epidemiological research, the curation of the data has been resource-intensive. There has been significant effort to develop and comply with data standards developed by SEER and the North American Association of Central Cancer Registries (NAACCR) (J. Swan et al., 1998; Wingo et al., 2003). The curation effort by cancer centers requires a staff of full time employees (FTE) to review and code a clinical case at an estimated pace of 5 cases per day per FTE (Kolender, 2009). In addition, central state cancer registries which collect case data and report on their statistics to state and federal agencies require an estimated 12.4 FTEs and an annual budget of $1 million (Chapman, Mulvihill, & Herrera, 2012; Tangka, Subramanian, Beebe, Trebino, & Michaud, 2010). Statistics derived from cancer registries have enabled stakeholders in healthcare to monitor the incidence of new cases, outcomes of patients, and identify opportunities for improvement.

Given the significant impact cancer care has on healthcare delivery and federal budgets, there has recently been a significant focus on the quality of cancer care. In the 1980s, many cancer centers worked to exempt themselves from the Medicare quality reporting efforts citing the complexity of cancer care. In 1999, significant quality improvement efforts in cancer care began with the landmark IOM publication Ensuring Quality Cancer Care (Hewitt & Simone,

1999), which defined ten cancer-specific quality recommendations. The IOM reports To Err is Human (Kohn, Corrigan, & Molla, 2000) and Crossing the Quality Chasm ("Crossing the Quality Chasm: A New Health System for the 21st Century," 2001) only added to this push. While the original exemptions led to a delay in the development of cancer-specific quality measures, the awareness generated by the IOM reports and new quality reporting standards have led to the ongoing effort to develop quality metrics (Spinks et al., 2011). A number of national organizations such as the National Comprehensive Cancer Network (NCCN), American Cancer Society (ACS), American Society of Clinical Oncology (ASCO), Quality Oncology Practice Initiative (QOPI), Commission on Cancer of the American College of Surgeons (CoC), College of American Pathologists (CAP), and American College of Radiology (ACR) are dedicated to quality improvement in cancer care by developing best-practice guidelines, engaging patients, and aggregating data for research (Winchester, Stewart, Phillips, & Ward, 2010). There have also been efforts by policymakers to derive more value from care by adjusting the incentive structures and developing new guidelines (Newcomer, 2012; Smith & Hillner, 2011). However, when done without engaging other stakeholders, value-driven changes can have unintended consequences (Jacobson, Earle, Price, & Newhouse, 2010). The recent efforts to generate cancer specific quality measures, guidelines, and patient engagement will enable the push towards high quality and value driven care.

While under the spotlight, advances in cancer research are making a positive impact on diagnostics, therapeutics, and clinical workflows. Approximately 20,000 cancer patients are enrolled in clinical trials each year by 14,000 investigators at more than 2,100 institutions and efforts are being made to streamline this process (Nass, Patlak, & Forum, 2013). In addition, there is a growing trend to expand the use of large simple and other pragmatic trials that focus on learning from real world care (Grossmann, Sanders, & English, 2013). Many of these trials,

spurred by whole genome and other next generation sequencing methods, are enabling scientists to understand the alterations that define cancer genomes (Meyerson, Gabriel, & Getz, 2010). This work is leading to progress in the use of genomic testing for cancer diagnostics (Meldrum, Doyle, & Tothill, 2011). In addition, this research has enabled the development of numerous FDA approved drugs that are targeted against specific molecular mechanisms instead of being cytotoxic to all cells (Vanneman & Dranoff, 2012). There have also been advances in the study of nanomaterials which could lead to advances in imaging and therapeutics (Barreto et al., 2011). New knowledge is being developed at such a rapid pace, that it will require near real-time quality metrics to evaluate the impact of the shifting cancer care landscape.

Traditional quality metrics have aimed to measure adherence to standard guideline based care. With new research on personalized diagnostics and therapeutics in addition to a greater role for patients in a shared decision making model, guidelines will become more complex. The clinical community will need robust quality metrics to represent the complexity of care while still promoting high quality, high value care.

2.3: Importance of Quality Metrics in Breast 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 et al., 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 et al., 2011). Breast cancer will continue to be one of the primary areas of care delivery, especially as the number of survivors continues to grow.

With the large population affected by breast cancer, efforts to promote awareness of the disease and fund research have grown. Despite the awareness raised by the "War on Cancer" by the National Cancer Act of 1971, breast cancer was still a very private disease through the 1970s and into the 1980s. From the founding of the Susan G. Komen foundation in 1982 through the National Breast Cancer Coalition in 1991, organizations and many individuals have raised awareness for the stories and issues faced by patients and their families. These organizations have led efforts to increase public and private investment in research, leading to federal research funding increasing from $81 million to more than $400 million in the 1990s. In addition, advocates have pushed for quality standards in clinical care, resulting in measures such as the Mammography Quality Standards Act in 1994. Finally, the numerous corporate partnerships that have been developed demonstrate how breast cancer advocacy has become mainstream (Braun, 2003; King, 2001; Sharf, 2001). This effort to raise awareness has led to significant research investment and focus on clinical quality.

With strong research efforts in breast cancer combined with a high level of patient engagement, knowledge of best practices is rapidly changing. Continuing advances in sequencing cancer genomes and understanding molecular pathways have enabled new diagnostics and therapeutics, with breast cancer tumors having been shown to have a median of 33 nonsynonymous mutations per tumor (E. D. Green & Guyer, 2011; Olopade, Grushko, Nanda, & Huo, 2008; Vogelstein et al., 2013). Advances have impacted diagnostics with commercial tests that range from detecting a single nucleotide variant (SNV) to full multiplex tests that utilize next generation sequencing methods to identify all major forms of gene alterations (C.L., Berger, & Pao, 2014; Domchek, Bradbury, Garber, Offit, & Robson, 2013). In addition, targeted

therapies are utilizing new knowledge to affect genetic alterations and signaling pathways with, for example, 10 FDA approved targeted therapies for breast cancer and 51 current NCI-sponsored clinical trials for breast cancer involving a targeted therapy ("Clinical Trials Search Results - National Cancer Institute," n.d., "Targeted Therapies for Breast Cancer Tutorial - National Cancer Institute," n.d.; Cortazar et al., 2012; M. J. Higgins & Baselga, 2011). In addition to genomic data, patient generated health data has the potential to inform clinical decision-making. As more breast cancer patients use mobile devices and become engaged with their care on digital platforms, patient generated data has the potential to inform the clinical decisions of the patient and their providers (Howie, Hirsch, Locklear, & Abernethy, 2014). Genomic and patient generated data have the potential to add to the body of knowledge around breast cancer, and continuous evaluation of quality metrics will be necessary to track changes in clinical practice.

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.

The initial diagnosis, staging, and tissue biomarker analysis of breast cancer is an important process in determining subsequent therapy. A suspected case of breast cancer is typically evaluated through breast imaging studies such as mammography, followed by a biopsy of the suspicious tissue which confirms the diagnosis (Smart, Hartmann, Beahrs, & Garfinkel, 1993). Clinical and pathologic tumor staging evaluates the tumor size, the lymph node status, and metastasis (TNM) using the American Joint Committee on Cancer (AJCC) (SB, DR, & CC, 2010) staging system. Histopathology represents the tumor cell origin while the grade represents the irregularity of the cancer cells. Standard tissue biomarker analysis includes evaluation of estrogen receptor (ER) expression, progesterone receptor (PR) expression, and human epidermal growth factor receptor 2 (HER2) expression or amplification. Receptor status informs both prognosis and opportunities for targeted therapies (Foulkes, Smith, & Reis-Filho, 2010).

Breast cancer has 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.

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) and the National Quality Measures for Breast Centers (NQMBC). Many professional organizations also have breast cancer quality components as part of their accreditation including The Quality Oncology Practice Initiative (QOPI), National Quality Forum (NQF) Breast Quality Measures, the College of American Pathology (CAP), the American College of Radiology (ACR), and the Commission on Cancer (CoC). 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 are 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.

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. An analysis of breast cancer specific quality metrics from the NQMBC, NAPBC, NQF, and QOPI revealed 71 unique quality metrics, some of which are overlapping across organizations. Of these metrics, 38 are clinical metrics, 30 are process metrics, and 3 are outcome metrics. Of the 71 metrics, 26 involved temporal features for the patterns of care. 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.

In this project, we study the following breast cancer quality metrics: 1) the rate of mastectomy vs. breast conserving surgery as the first surgical event; 2) the rate of re-excision after breast conserving surgery; 3) the rate of radiation therapy after breast conserving surgery; and 4/5) the rate of pre-operative or post-operative chemotherapy usage (Figure 1). The rate of breast conserving surgery compared to mastectomy as the first surgical event provides a window into the downstream effects based on the potential resulting treatment paths. Re-excision of cancerous tissue occurs after a breast conserving surgery results in inadequate surgical margins, and can lead to psychological, physical, and economic stress for patients. Furthermore, it can delay adjuvant therapies and even lead to the need for full mastectomy. As a result, the re-excision rate after a breast conserving surgery has been an important indicator of surgical quality (McCahill et al., 2012). Radiation therapy can reduce local recurrence rates when delivered after a breast conserving surgery. As a result, it has been a valuable quality metric to track the clinical protocols in place at cancer centers (Korah, Sener, & Tripathy, 2012; Vinh-Hung & Verschraegen, 2004). Finally, we assess two chemotherapy usage quality metrics.

First, patients with hormone-receptor negative tumors have been shown to respond to neo-adjuvant chemotherapy, and thus we assess the rates of neo-adjuvant chemotherapy (Carey et al., 2007; Liedtke et al., 2008; von Minckwitz & Martin, 2012). The second assessment focuses on provider decision-making in response to the results of the Oncotype DX diagnostic test. The Oncotype DX test can predict recurrence in lymph node negative, estrogen receptor positive patients, and thus would affect the decision to administer chemotherapy (Asad et al., 2008; Dabbs et al., 2011; Flanagan, Dabbs, Brufsky, Beriwal, & Bhargava, 2008). This broad set of quality metrics will require patient treatment event sequences to be represented in many forms based on the specific quality metric being calculated.

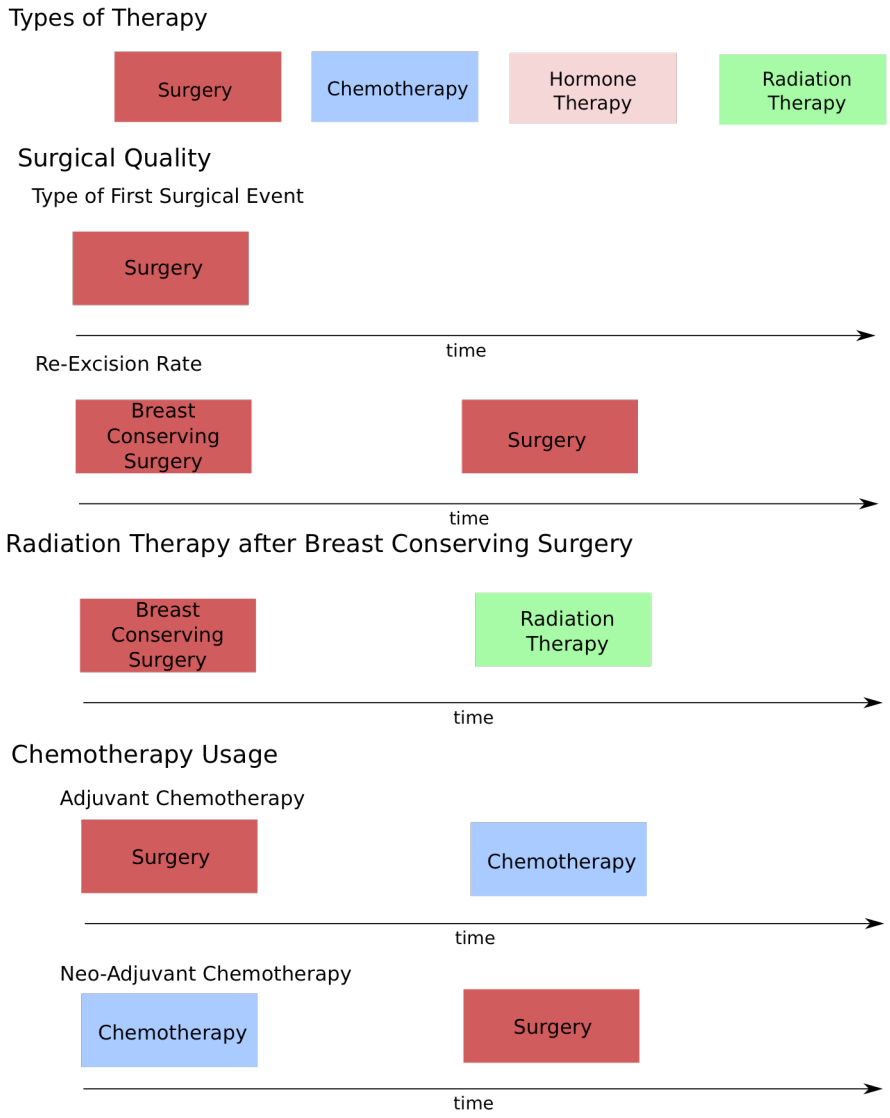

Figure 1: Patterns of breast cancer quality metrics

Top: 4 primary modalities of breast cancer therapy - surgery, chemotherapy, hormone therapy, radiation therapy

Bottom: Represents the quality metrics assessed in this work and a simple example of the event pattern that encodes the quality metric.

2.4: Challenges and Opportunities for Quality Metrics in Breast Cancer Care

Generating quality metrics from clinical 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 assess for quality. 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 quality metric or clinical question. NAACCR, SEER, the NCI Thesaurus, and others work to meet this challenge by developing hierarchical ontologies and dictionaries to represent clinical data. 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.

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 2 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 2 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 sequence in order to calculate clinical quality metrics.

Finally, delivering quality metrics to stakeholders effectively in a continuous fashion is a major challenge. The development of these metrics 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. 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 the system to learn from every clinical encounter.

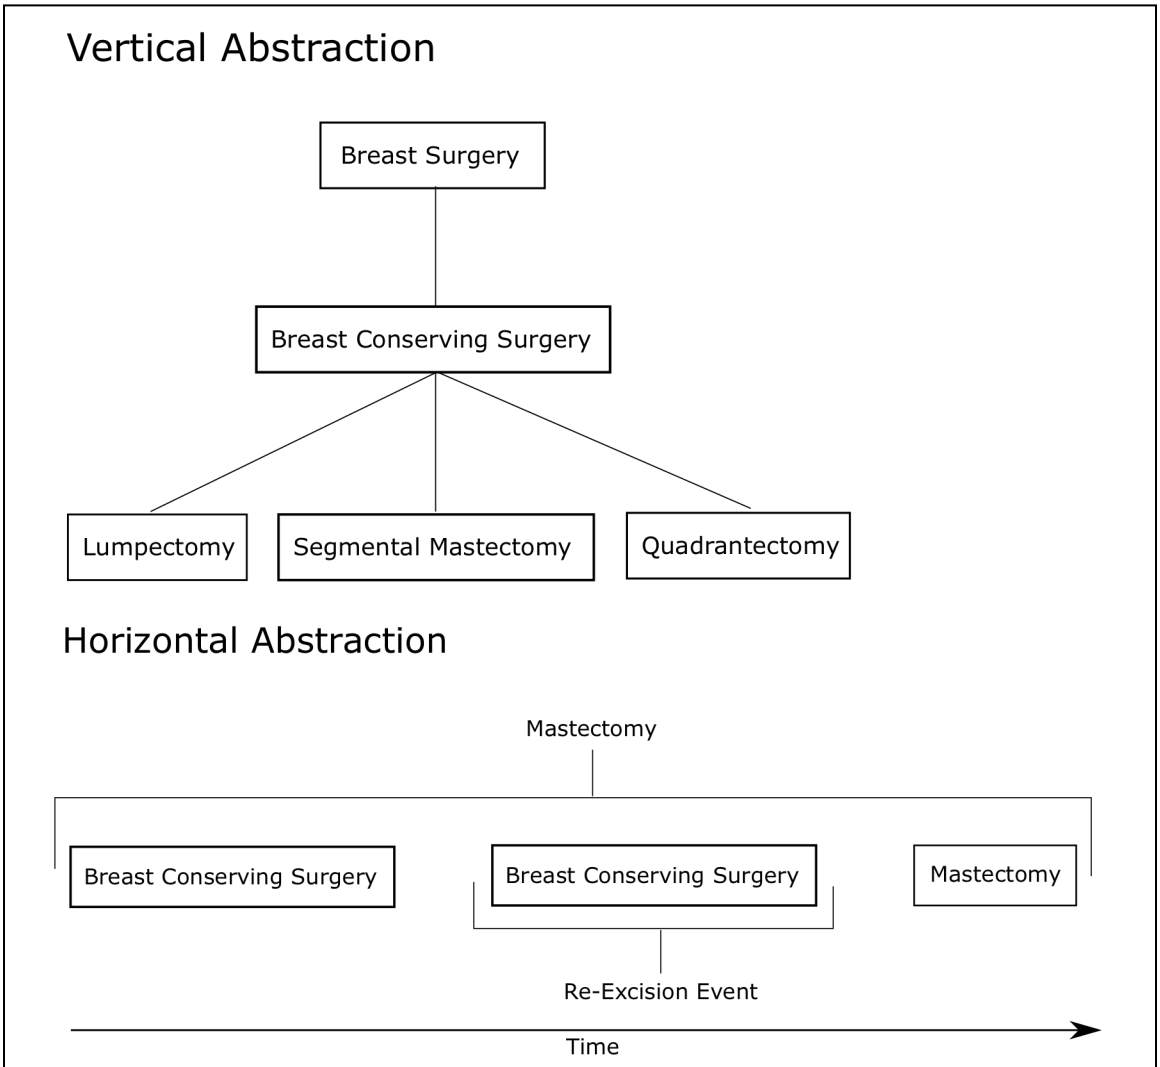

Figure 2: Example of vertical and horizontal abstraction for linking clinical data with quality metrics

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

Tracking clinical care in near real time requires multiple subtasks to convert raw clinical data to reportable quality metrics (Section 3.1). 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. Visualization methods have allowed clinical experts to identify clinical patterns from the abstracted information. In addition, event sequence mining methods have helped automate the pattern identification process (Section 3.2). Temporal abstraction frameworks consist of basic components which define how to encode the clinical knowledge and the inputs and outputs of various data transformation subtasks (Section 3.3). This study will aim to utilize vertical and horizontal abstraction methods to generate clinical quality metrics for breast cancer treatment by using cancer registry and billing code data (Section 3.4).

3.1: Imperative for Tracking Clinical Care Patterns Over Time

Tracking clinical care metrics in near real-time 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 about quality. 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 generating clinical quality metrics. 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 quality metrics (Figure 3).

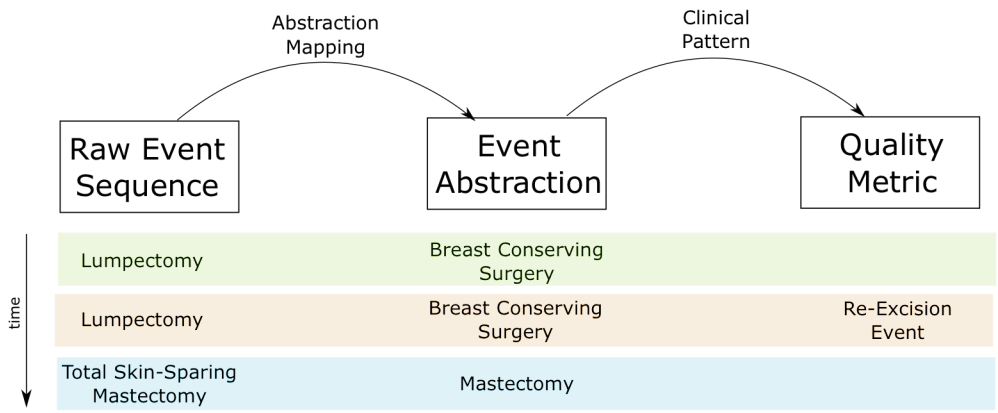

Figure 3: Converting raw event sequences to information for quality metrics

Generating complex clinical quality metrics 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: Methods for Tracking Clinical Care Patterns

Prior methods for tracking clinical care patterns have included heuristic, knowledge based, or 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.

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 & Musen, 1996; Yuval Shahar, 2013). 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, Fagan, & Sheiner, 1991; M. G. Kahn, Fagan, & Tu, 1991; M. G. Kahn, 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.

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.

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). 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.

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.

Generating quality metrics for clinical care in near real-time 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 quality metrics. 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 clinical quality metrics for breast cancer care that require event pattern sequences.

3.3: Temporal Abstraction of Clinical Care Patterns

Formal models describing the input and output, relations between entities, and context and domain specificity for the abstraction process are an important part of operationalizing and generalizing the temporal abstraction framework. The KBTA framework and other temporal abstraction systems have delineated this set of definitions. The KBTA framework consists of five

parallel subtasks: temporal context restriction, vertical temporal inference, horizontal temporal inference, temporal interpolation, and temporal pattern matching. This framework enables the interpretation of raw clinical data via a higher-level descriptive terminology. These features are vital in developing and visualizing clinical quality metrics.

The transition from a large set of raw clinical data to an easily interpretable quality metric requires a formal understanding of the input and output entities and how they relate. Main ontological entities include clinical parameters (i.e., laboratory values and other continuous variables), events (i.e., surgery and other discrete events), abstraction goals (i.e., complex, context-sensitive objective), clinical patterns (i.e., patterns of clinical parameters with time and value constraints), and interpretation contexts (i.e., clinical context for abstraction process). The input to the system consists of time stamped clinical parameters and events along with the set of abstraction goals. The output is a temporally dependent, context-specific parameter that is at the same or higher abstraction level depending on, in our case, the desired quality metric. The logical proposition of an abstraction includes the value of a parameter in a specific clinical context in a specific time interval.

There are four defined groups of domain-based relations: structural knowledge, classification knowledge, temporal-semantic knowledge, and temporal dynamic knowledge. Structural knowledge denotes relationships between different clinical entities, such as is-a, and part-of relations. Classification knowledge demonstrates how clinical entities are grouped based on, for example, the value of a clinical parameter or the temporal pattern of a clinical event series. Temporal semantic knowledge indicates how the time intervals for different clinical entities are abstracted and interpreted based on the entity and clinical context. Finally, temporal dynamic knowledge guides how clinical parameters and events persist and can be represented at times when their value is not measured. These four types of knowledge enable

the five subtasks of a formal abstraction methodology.

The first of five primary subtasks is temporal context restriction that identifies the clinical context under which to limit the scope of inference. The interpretation context for clinical data serves as a temporal frame of reference and can be prospective, retrospective, or for the current time interval. The second subtask is vertical temporal inference that identifies higher-level concepts that occur in the similar time interval. Horizontal temporal inference is the third subtask and represents inference across multiple clinical entities from different time intervals using context and interval-based logic. The fourth subtask is temporal interpolation that focuses on filling the gaps between disjoint clinical entities using context-specific truth persistence functions to understand how clinical parameters change over time. The final subtask is temporal pattern matching which identifies a clinical pattern over a disjoint set of intervals and creates a new entity and interval based on predefined patterns.

Temporal abstraction methodologies have been primarily implemented to tackle issues facing clinical decision support (CDS). We believe these principles can also be used to surmount the challenge of generating near real-time clinical quality metrics. Ontological entities such as clinical events, parameters, patterns, abstraction goals, and interpretation contexts are important for verifying whether the desired patterns of care occur in a specific context. The four categories of knowledge (structural, classification, temporal semantic, temporal dynamic) can be used with the input (events, parameters, abstraction goals) to run the primary subtasks (temporal context restriction, vertical temporal inference, horizontal temporal inference, temporal interpolation, temporal pattern matching). This process can allow for the identification of patterns from clinical data and linkage with quality metric patterns that exist at different levels of abstraction.

3.4: Hypothesis and Aims

We believe that temporal abstraction principles can be utilized to generate near real-time clinical quality metrics. The goal of this work is to develop a method to continuously monitor clinical quality metrics that have a temporal component for breast cancer treatment in near real time. Our objective is to develop a scalable framework to identify and visualize patterns of care at multiple levels of abstraction and use it to generate clinical quality metrics. We hypothesize that a framework that consists of vertical and horizontal abstraction methods can help generate these quality metrics.

Aim 1: Develop a scalable framework to identify and visualize patterns of care at multiple levels of abstraction

Aim 2: Evaluate framework using registry data to characterize breast cancer patterns of care and evaluate quality metrics

Aim 3: Evaluate potential use of administrative data for near-real time metrics

Chapter 4: The Pathfinder Framework

This chapter describes the Pathfinder method and its application to determining breast cancer quality metrics. The data for this study is derived from the Vanderbilt University Medical Center (VUMC) cancer registry system and administrative data (Section 4.1). The Pathfinder method consists of six sub-tasks including data extraction, data standardization, vertical abstraction, horizontal abstraction, quality metric querying, and generating/visualizing quality metrics (Section 4.2). We extracted cancer registry treatment events and CPT codes, mapped them to the NCI Thesaurus, and implemented our vertical and horizontal abstraction methods (Section 4.3). We next calculated our specified breast cancer clinical quality metrics on the abstracted cancer registry event data (Section 4.4). Finally, we compared the CPT codes against the cancer registry events and ran the surgical quality metric on the administrative data (Section 4.5).

4.1: Clinical Setting and Patient Data Sources

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

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. A team of specially trained nurse registrars maintains the database using the METRIQ cancer registry data management system from Elekta. The registry nurses use 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 from the VUMC medical record system and external information sources. 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.

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.

4.2: Pathfinder Framework for Generating Clinical Quality Metrics

The general methodology for this study consists of six major subtasks: data extraction, data standardization, vertical abstraction, horizontal abstraction, quality metric filtering, and generation/visualization of quality metrics (Figure 4). Data extraction involves the collection of the raw, time-stamped 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 quality metric filtering subtask pulls the event sequences in the context of the desired quality metric by utilizing the clinical pattern of the quality metric and identifying the proper level of vertical and horizontal abstraction to use. Finally, the quality metric generation and visualization subtask involves using the final set of treatment events to calculate and visualize the desired information regarding the metric. This series of six subtasks converts raw data to a form that can generate a variety of clinical quality metrics.

This methodology was implemented on a local secure server. The raw data to be read for the data extraction subtask was pulled from the RD. The data standardization, vertical abstraction, and horizontal abstraction subtasks are implemented in a perl environment ("The Perl Programming Language - www.perl.org," n.d.) and the treatment event sequences and their abstractions were stored on the Netezza appliance. The visualizations were generated using the Google Charts JavaScript platform Visualizations will utilize the Google Charts JavaScript

platform (Google, 2014).

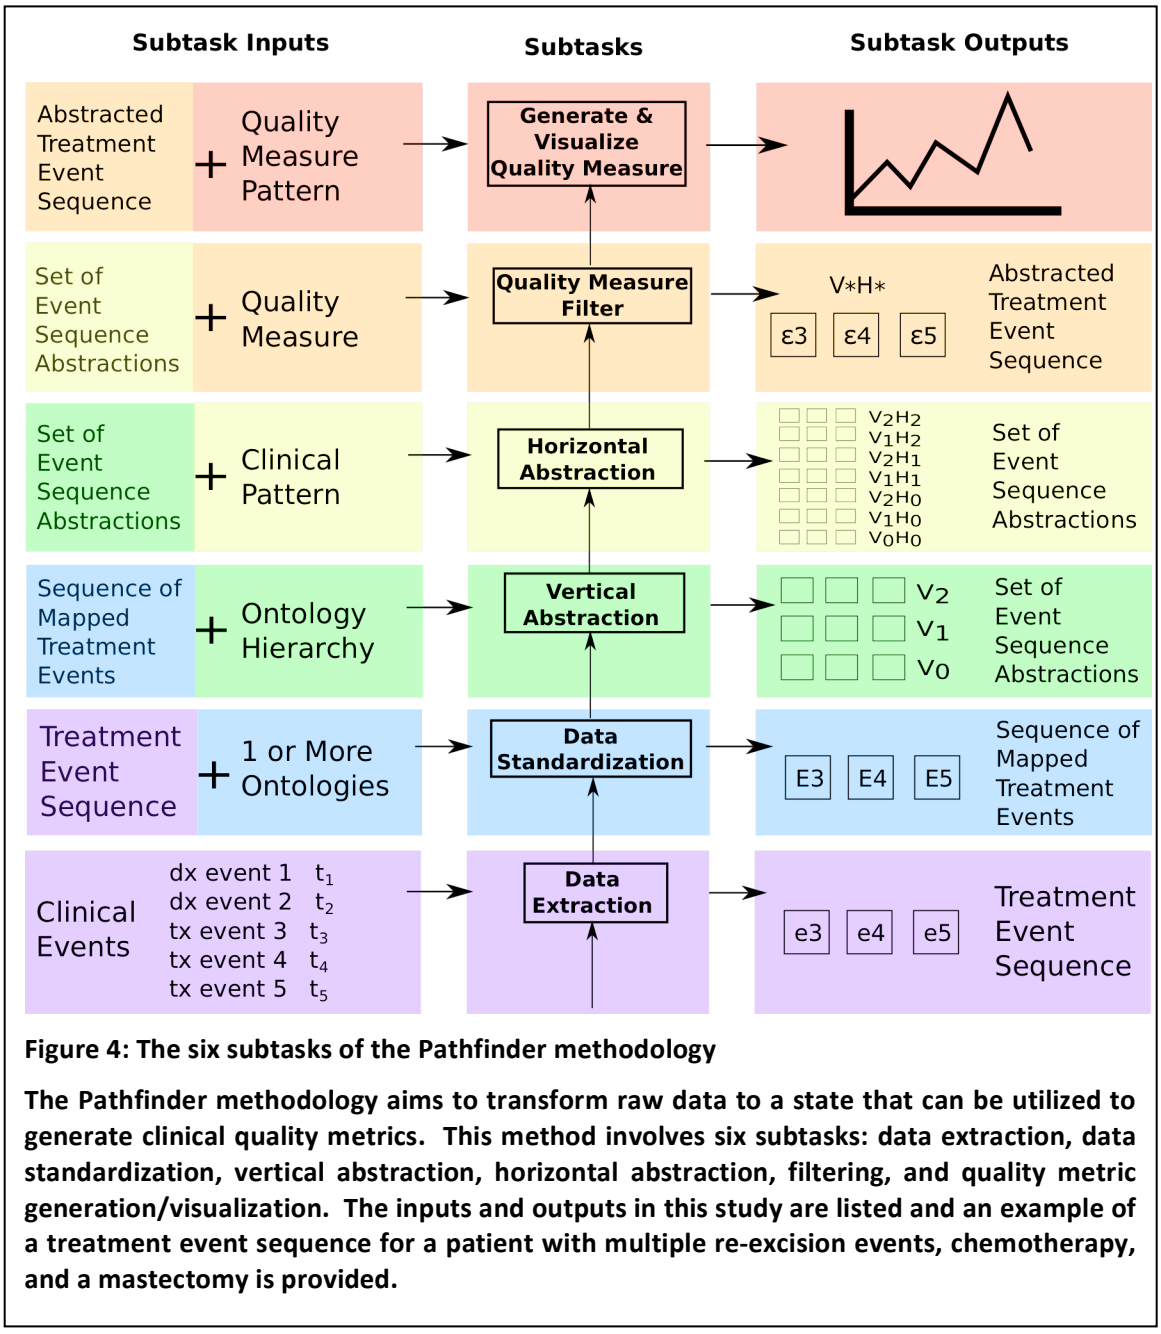

Figure 4: The six subtasks of the Pathfinder methodology

The Pathfinder methodology aims to transform raw data to a state that can be utilized to generate clinical quality metrics. This method involves six subtasks: data extraction, data standardization, vertical abstraction, horizontal abstraction, filtering, and quality metric generation/visualization. The inputs and outputs in this study are listed and an example of a treatment event sequence for a patient with multiple re-excision events, chemotherapy, and a mastectomy is provided.

4.3: Extraction and Representation of Patient Data

The data extraction subtask consisted of selecting treatment records for adult women diagnosed with breast cancer between 2000 and 2012 from the cancer registry. Treatment event data on chemotherapy, immunotherapy, surgery, hormone therapy, and radiation therapy were extracted from the cancer registry while surgical CPT codes were extracted from the administrative data. We also converted experimental therapies, listed in the cancer registry as event type "O", to their therapy type (i.e., chemotherapy). Women who had bilateral breast cancer, had a recurrence of breast cancer, or were on their second course of therapy were excluded from this study. Additionally, women were only included if they were on a curative course of therapy (Stage I-III breast cancer) and had at least one surgical, chemotherapeutic, immunotherapeutic, radiation, or hormone therapy event. Only the treatment paths from the first occurrence of disease and first course of care were included. Of the selected patients, 10% were randomly set aside as a holdout set for future use. We characterized the patient cohort by the patients' age at diagnosis, race, primary cancer site, histology, and clinical stage.

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 5). 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 and CPT 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.

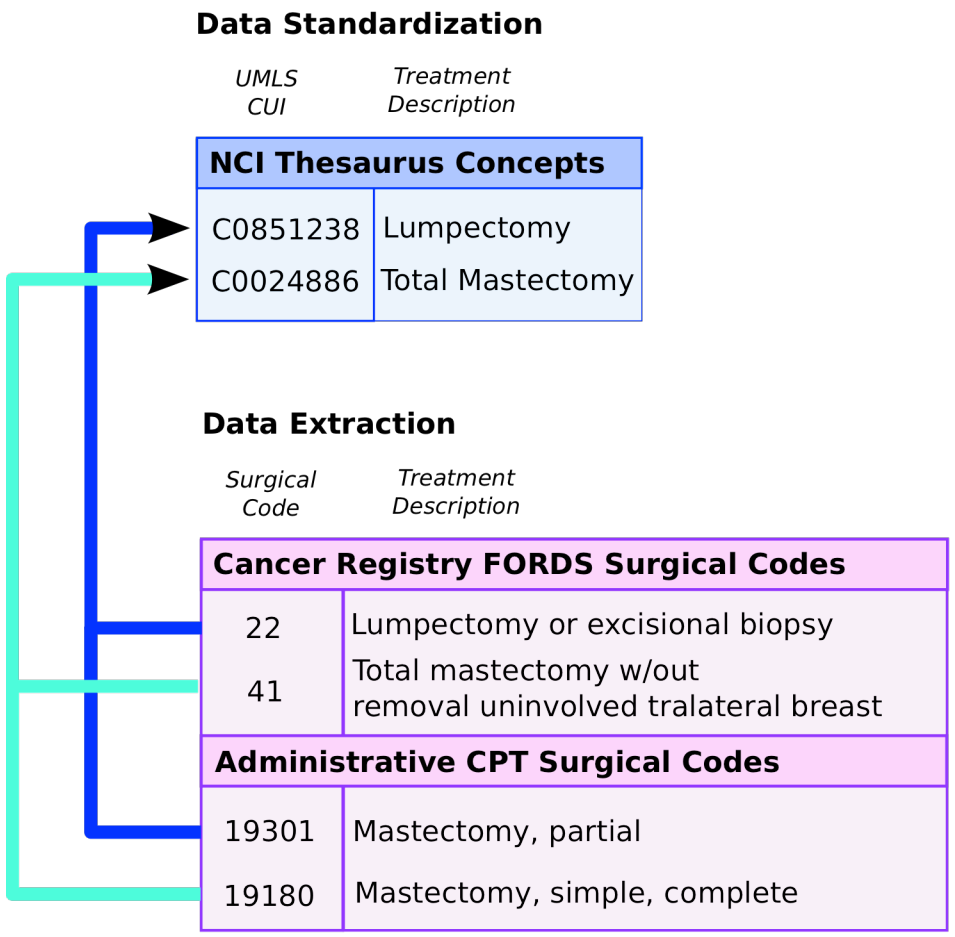

Figure 5: Example of the Data Extraction and Standardization subtasks

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 standardization step.

The vertical abstraction subtask involved the parsing of the ontological hierarchy as well as consolidating similar events that occur simultaneously. A subset of the NCI Thesaurus hierarchy generated the vertical abstractions for surgical treatment events (Figure 6). We used three levels of vertical abstraction to represent the specific mastectomy and breast conserving 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.

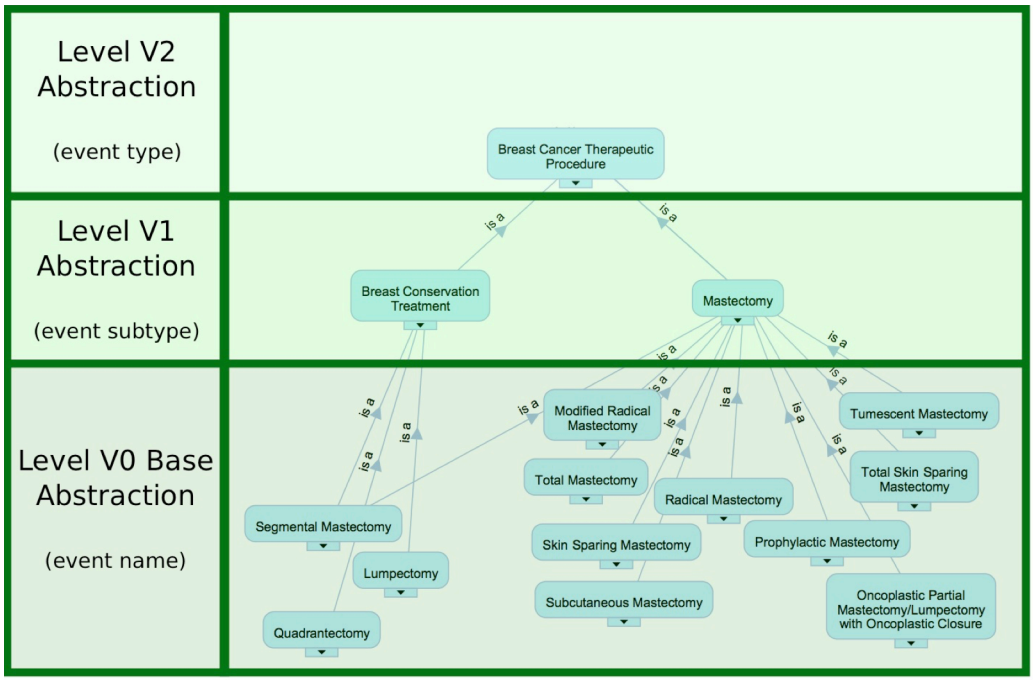

Figure 6: Example of Vertical Abstraction subtask for surgical breast cancer treatment

The NCI Thesaurus ontology tree for breast cancer therapeutic procedure is displayed. 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 7). 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.

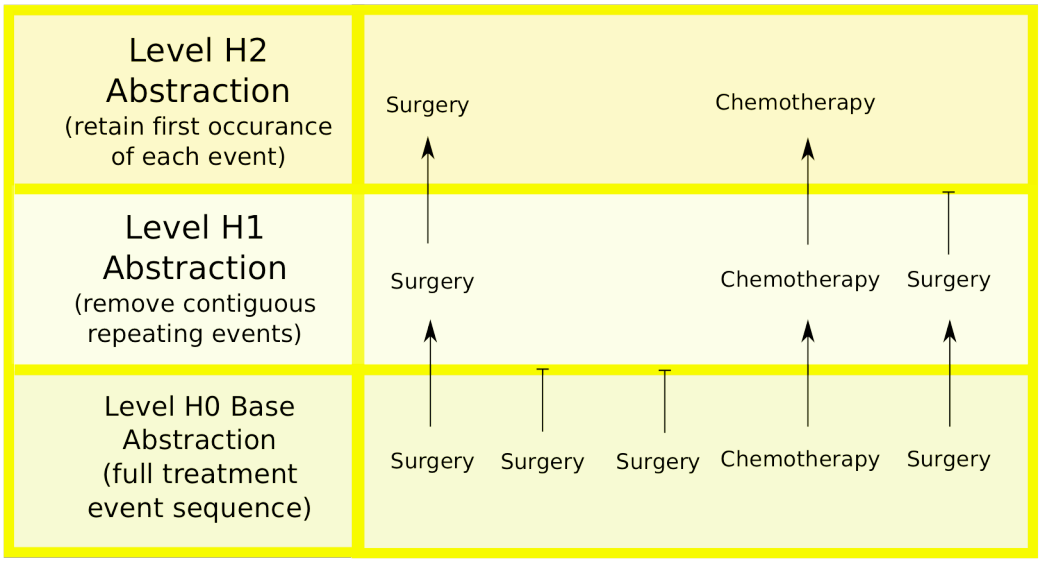

Figure 7: Example of Horizontal Abstraction subtask for treatment event sequences

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. We also focused on visualizing the treatment event sequences by generating a set of Sankey diagrams(Schmidt, 2008) using Google Charts. The Sankey visualization technique is ideal for representing the magnitude of flow between sequential nodes that represent treatment events. We represented each treatment event as a node in the Sankey diagram to show the number of patients that were represented by the various event sequences. We developed Sankey diagrams at the V1H0, V1H1, V1H2, V2H0, V2H1, and V2H2 levels of abstraction to provide a visual representation of the abstraction subtasks and quantified the changes in number of nodes and links. Finally, we used the SPADE event sequence mining method in the aRules package (Hahsler & Chelluboina, 2011; Zaki, 2001) in the R statistical computing environment (R Core Team, 2013). We identified the set of event sequence patterns that had a support value greater than 0.002 from the set of treatment event sequences of stage I-III breast cancer patients without using specific time ranges. We identified these frequent event patterns for abstraction levels V0H0, V1H0, V1H1, V1H2, V2H0, V2H1, and V2H2. We ranked the event sequence patterns by their support value and calculated the number of treatment event sequences and patients the most frequent event pattern was present in. Further, we calculated the cumulative treatment event sequence and patient coverage by the frequent event patterns for additional patterns. This analysis will allow for a comparison on how well frequent event patterns abstract the complexity of treatment event sequences across the population. These three characterization steps enable the

assessment of the results through the prism of the treatment event sequences.

4.4: Generating Quality Metrics from Abstracted Cancer Registry Data

The quality metric filtering subtask involves the use of a clinical pattern that represents the quality metric in order to identify the optimal vertical and horizontal abstraction levels. In this study, we focus on re-excision rates, rate of radiation therapy after breast conserving surgery, and rates of chemotherapy usage.

Querying for the re-excision rate quality metric requires the definition of a clinical pattern that can identify a re-excision event. A re-excision event occurs anytime there is a breast conserving surgery after an initial breast conserving surgery event. This clinical pattern was queried against surgical events at abstraction level V1H0 for stage I-III patients (Figure 8). We also produced a Sankey diagram for the re-excision rate quality metric in order to provide a visual representation of the patients' treatment event sequences. The second metric we assess is the rate of radiation therapy after breast conserving surgery for stage I-III patients. The clinical pattern is a radiation treatment event that occurs after one or more breast conserving surgeries. This metric utilizes level V1H0 abstraction for surgical events and level V2H1 abstraction for radiation treatment events. All surgical procedures must be breast conserving, and never advance to a total mastectomy. On the other hand, only the first radiation therapy event of a sequence is required to determine if it took place after a breast conserving surgery.

The final set of quality metrics focus on chemotherapy usage. First, we assessed the rate of neo-adjuvant chemotherapy for stage II-III patients with hormone-receptor negative tumors. The clinical pattern is a chemotherapy event that takes place before a surgery event in a patient who has no hormone therapy events. As a result, a level V2H2 abstraction that shows the order of events is used. The final metric focuses on the rate of chemotherapy hormone

receptor positive, HER2 receptor negative, lymph node negative, stage I-II breast cancer patients. This metric also utilizes the level V2H2 abstraction to identify if a chemotherapy event occurred or not. In addition, we used the presence of hormone therapy and lack of immunotherapy as a predictor of receptor status. Clinical patterns representing quality metrics can enable the proper usage of vertical and horizontal abstractions.

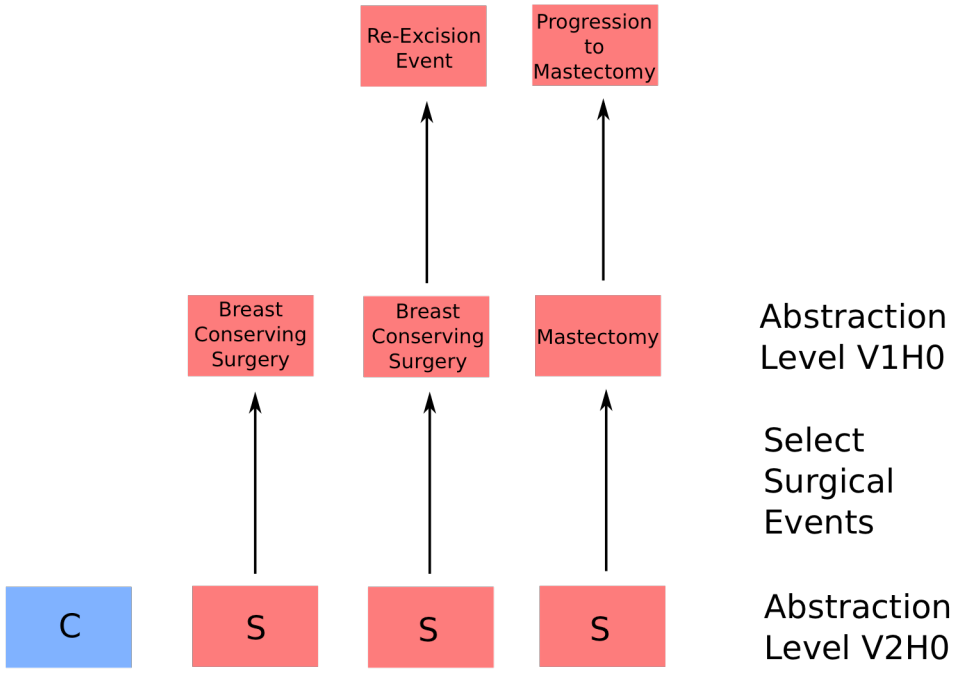

Figure 8: Querying for the re-excision rate quality metric

The process of filtering for and generating quality metrics involves the application of a quality metric pattern to a set of abstracted treatment event sequences. For the re-excision quality metric, the surgical events are selected at the V2H0 level of abstraction and the first type of surgical event is assessed at the V1H0 level of abstraction. If it is a breast conserving surgery, the following surgeries are interpreted as either a re-excision event or progression to mastectomy (depending on the type of surgery at the V1 level).

The final subtask of generating and visualizing the quality metrics involve defining the format for presenting the results. We calculated the rate of the metric-specific clinical patterns' presence in the designated abstraction level for the patient cohort for each year of diagnosis from 2000 to 2012. A 95% confidence interval was generated for all proportions for each year

of diagnosis using the exact method. We also used the Chi-square test for trends to determine if there was a statistically significant difference between different quality metric rates over time.

4.5: Characterizing Abstracted CPT Data and Generating Quality Metrics

Surgical CPT codes from the VUMC clinical information system were extracted and used to assess the potential for using non-registry data to measure clinical quality in near real time.

Our first objective was to analyze the number of CPT codes that were missing or misaligned with surgical events in the cancer registry. We first extracted any breast surgical CPT code (CPT=19***) that occurred within one day of a cancer registry surgical event. We counted the number of times that a CPT code was present for surgical events in the cancer registry that took place or did not take place at VUMC. We repeated this analysis with only CPT codes that specifically represented a mastectomy or breast conserving surgery. We also used the specific mastectomy and breast conserving surgery codes to measure the fraction of codes that were missing in various positions of the surgical event sequence. We calculated the number of CPT codes that were improperly aligned with the type of surgery represented in the cancer registry. Finally, we counted the number of patients with at least one missing mastectomy or breast conserving surgery CPT code in their surgical event sequence.

Our second objective was to compare the surgical event sequences for our patient cohort. We extracted all mastectomy and breast conserving surgery CPT codes for stage I-III breast cancer patients and used the codes that occurred within two years of the patients' first CPT code to generate our surgical event sequences. We re-calculated the re-excision quality metric based on CPT data. For this, we selected all mastectomy and breast conserving surgery CPT codes for stage I-III breast cancer patients. As with the cancer registry analysis, we used the V1H0 level abstraction. The rates of breast conserving surgery and mastectomy as well

as the rate of re-excision was calculated using abstracted CPT codes and compared to the similar analysis conducted for the cancer registry. We used the Chi-square test for trends to compare the CPT rates against the cancer registry rates. We also directly compared surgical event sequences for stage I-III breast cancer patients that originated from CPT and the cancer registry. We measured the fraction of surgical event sequences that were equivalent at abstraction levels V0H0, V1H0, V1H1, V1H2, and V2H0. Because we are only assessing surgical events, abstraction levels V2H1 and V2H2 would over-consolidate the sequence and were not assessed.

Our final objective was to compare the rate at which frequent event patterns occur in CPT and cancer registry derived surgical event sequences. We used the same surgical event sequences extracted for the second objective of this section. We implemented the same event sequence mining method used previously but instead used the surgical event sequences at the V1H0 abstraction level. For each frequently occurring event pattern, we calculated the number of patients that the pattern occurred in from the CPT and cancer registry derived surgical sequences. We compared the number of patients with or without the pattern using the Chi-square test. This analysis provides a view into the opportunities and challenges for using administrative data for quality metrics in real time.

Chapter 5: Results of the Pathfinder Methodology

This chapter describes the results of the Pathfinder method implementation and evaluation. We extracted the treatment event records for 2679 patients that were diagnosed with breast cancer between 2000-2012 and met our inclusions criteria from the VUMC cancer registry (Section 5.1). We next characterized the results of the vertical and horizontal abstraction process and demonstrated the 12-fold reduction in number of unique treatment event sequences from raw abstraction to level V2H2 (Section 5.2). Section 5.3 reviews the re-excision rate, radiation after breast conserving surgery, and chemotherapy usage quality metrics resulting from the cancer registry data analysis. We then characterized the CPT data by measuring the amount of missing data and their alignment with the cancer registry treatment event data (Section 5.4). Finally, section 5.5 presents the re-excision quality metric resulting from the CPT analysis and compares its performance to the cancer registry data.

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

We originally extracted the records of 4378 breast cancer patients diagnosed between 2000-2012 from the VUMC cancer registry that had undergone some diagnostic or treatment event. We set aside 453 records for our holdout set. Out of the remaining 3925 records, 2679 had fully reported NAACCR records, met our inclusion criteria, and were thus utilized for this study (Figure 9). These 2679 cases had 8023 treatment events and 72 experimental treatment events recorded in the tumor registry. Of the 72 experimental treatment events, 68 were classified as chemotherapy events and 4 as hormone therapy events. The demographic characteristics for this population are presented in Table 1. The patient population selected appears representative of the national breast cancer population, although seems to

underrepresent the African-American population (DeSantis, Ma, Bryan, & Jemal, 2014). Finally, Figure 10 illustrates the distribution of stage I-III breast cancer patients in our cohort across their year of diagnosis. The number of patients diagnosed per year at VUMC and included in the cancer registry has grown from 73 in 2000 to 226 in 2012.

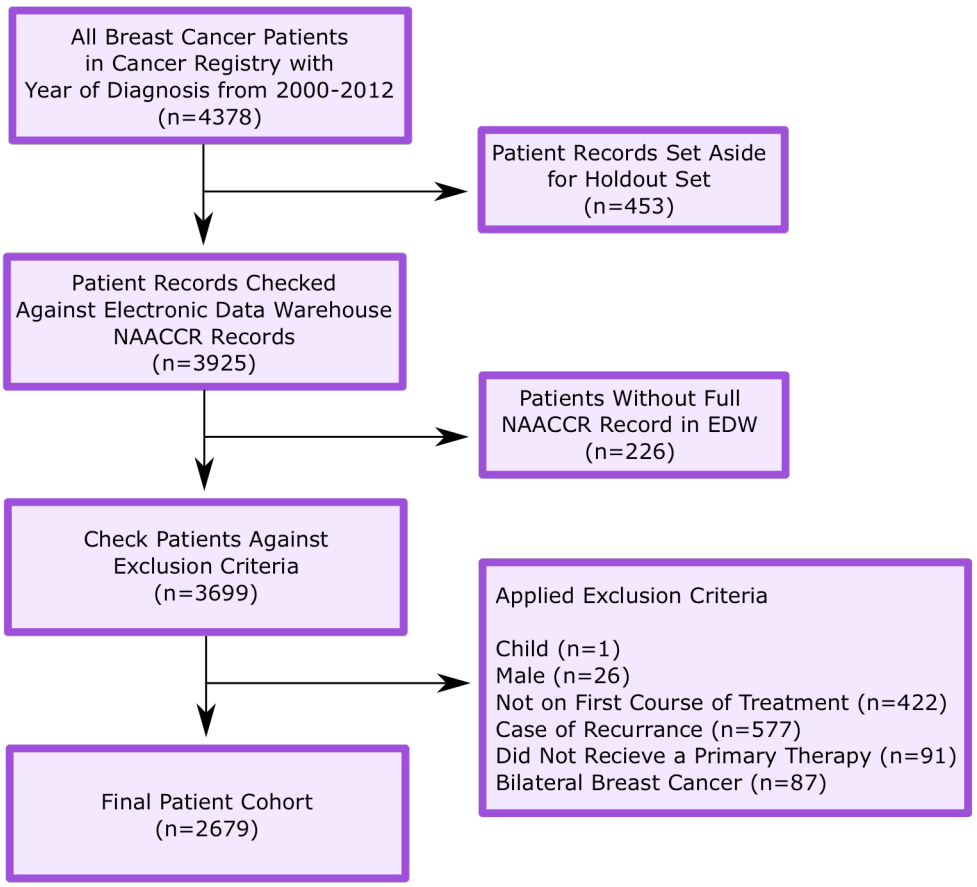

Figure 9: Breast cancer patient cohort selection process using cancer registry data

Patient cohort selection involving the application of exclusion criteria, confirming the presence of corresponding NAACCR records in the electronic data warehouse (EDW), and the setting aside of the holdout set. The various exclusion criteria did overlap leading to the removal of 1020 patients (instead of 1204).

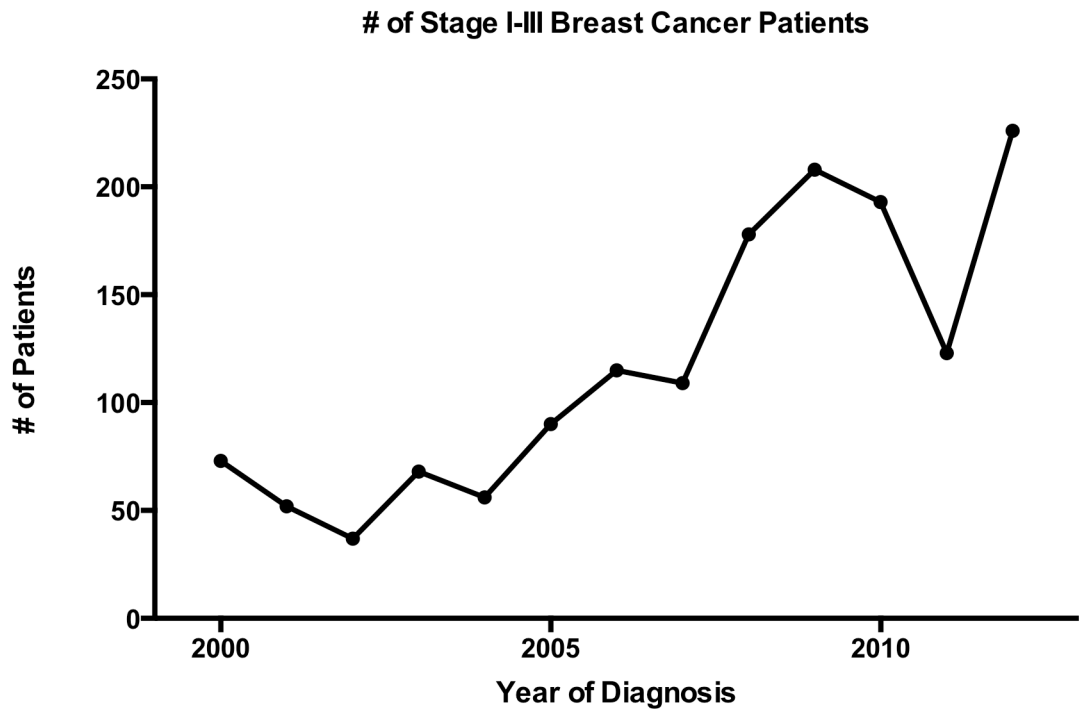

Figure 10: Number of stage I-III breast cancer patients diagnosed at Vanderbilt and included in the cancer registry from 2000-2012

The number of stage I-III breast cancer patients diagnosed per year and included in the cancer registry has grown from 73 in 2000 to 226 in 2012.

Table 1: Demographic characteristics of breast cancer patient cohort

Characteristic	Median	1st Quartile	3rd Quartile	Min	Max
Age at diagnosis (N=2679)	55	47	64	21	99

Characteristic	Category	Count
Race (top 3)	**Total**	**2679**
	White	2301 (85.9%)
	Black	304 (11.3%)
	Other Asian, including Asian or Oriental, NOS	22 (0.8%)
Primary Site (all)	**Total**	**2679**
	C500, Nipple	29 (1.1%)
	C501, Central portion of breast (subareolar)	159 (5.9%)
	C502, Upper inner quadrant	248 (9.3%)
	C503, Lower inner quadrant	167 (6.2%)
	C504, Upper outer quadrant	1038 (38.7%)
	C505, Lower outer quadrant	176 (6.6%)
	C506, Axillary tail	8 (0.3%)
	C508, Overlapping lesion of breast	541 (20.2%)
	C509, Not otherwise specified (NOS)	313 (11.7%)
Histology (top 5)	**Total**	**2679**
	8500/3, Infiltrating duct carcinoma, NOS	1326 (49.5%)
	8010/3, Carcinoma, NOS	339 (12.7%)
	8500/2, Intra-ductal carcinoma, non-infiltrating, NOS	186 (6.9%)
	8520/3, Lobular carcinoma, NOS	183 (6.8%)
	8522/3, Infiltrating duct and lobular carcinoma	123 (4.6%)
Clinical stage (all)	**Total**	**2679**
	0	379 (14.1%)
	1	853 (31.8%)
	2	503 (18.8%)
	3	171 (6.4%)
	4	122 (4.6%)
	Unknown/Not documented	651 (24.3%)

Demographic characteristics are provided for the 2679 patient records in the breast cancer patient cohort. Characteristics derived from the cancer registry include age, race, primary tumor site, histology, and clinical stage.

5.2: Data Standardization and Abstraction from Cancer Registry Data

We carried out the data standardization subtask for cancer registry treatment events by utilizing our mapping of FORDS surgical codes and NCI Thesaurus concepts. Treatment events were placed in temporal order for each patient to develop the set of treatment event sequences. We next carried out the vertical and horizontal abstraction subtasks using our ontology hierarchy.

We next characterized the treatment event sequences for the 1528 stage I-III breast cancer patients. We first counted the number of unique treatment event sequences for patients from 2000-2012 at varying levels of vertical abstraction for surgical events and horizontal abstraction for all events. The vertical and horizontal abstraction process led to a 12-fold reduction in the number of unique treatment strings from the least abstracted (V0H0) to most abstracted (V2H2) representation (Table 2). We also counted the number of most frequently occurring treatment event sequences that represented 90% of the patient population. This demonstrated how higher levels of abstraction consolidated the number of event sequences reduced the amount of variability. While 62.7% of event sequences were necessary to represent 90% of the population at the V0H0 level of abstraction, only 35.5% of the event sequences were necessary at the V2H2 level. This trend persisted as most treatment event sequences were consolidated through the abstraction process (Figure 11). This represents the simplification of treatment pathways leading to variable sizing of patient cohorts for study and analysis. The data standardization subtask (raw data to level V0H0) reduced the number of sequences due the post-coordination of cancer registry event names (i.e., concept for mastectomy with or without reconstruction). Finally, we graphed the number of patients per treatment event sequence at the V2H1 level of abstraction (Figure 12). The treatment event sequences exhibit an exponential distribution when ranked by the number of patients represented by the sequence. The top 10 most frequent treatment event sequences begin with a surgery while the next nine event sequences begin with chemotherapy.

Table 2: Number of unique treatment event sequences for each abstraction level

Level of abstraction	# of unique treatment event sequences	# of unique treatment event sequences to cover 90% of population
Raw data from data extraction	632	480 (75.8%)
Level V0H0 abstraction	391	245 (62.7%)
Level V1H0 abstraction	265	126 (47.5%)
Level V1H1 abstraction	154	54 (35.1%)
Level V1H2 abstraction	73	28 (38.4%)
Level V2H0 abstraction	205	79 (38.5%)
Level V2H1 abstraction	97	26 (26.8%)
Level V2H2 abstraction	45	16 (35.5%)

The number of unique treatment event sequences for each level of abstraction. There is a 12-fold reduction in number of sequences from the raw event sequence to the level V2H2 abstraction. The number of most frequently occurring event sequences needed to cover 90% of the patient population is also presented.

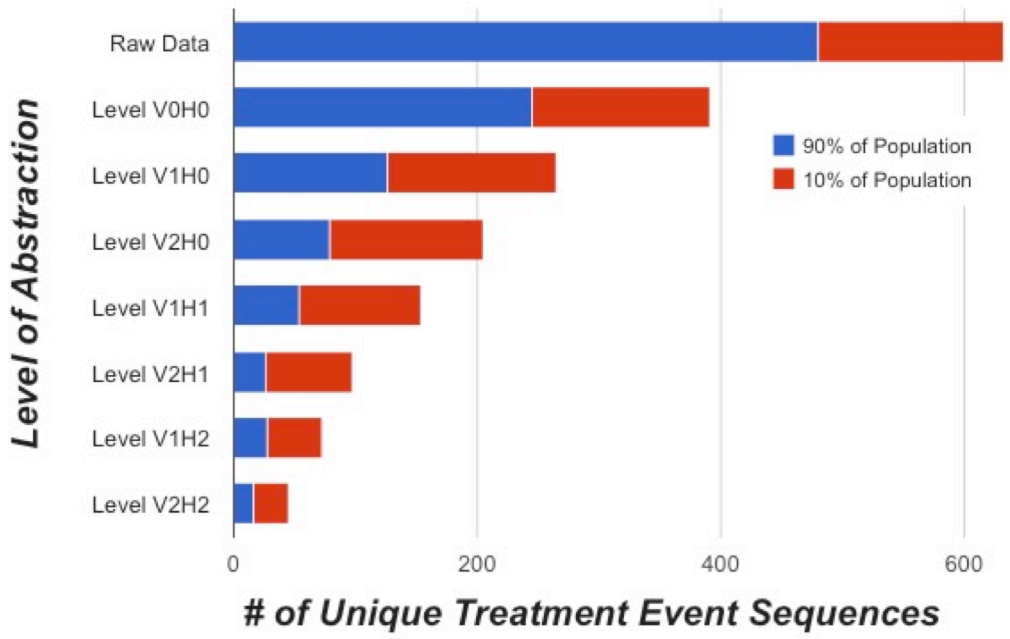

Figure 11: Fraction of treatment event sequences that are consolidated through abstraction

Most treatment event sequences are consolidated through the abstraction process. From ther raw/V0H0 to V2H2 level of abstraction, the fraction of treatment event sequences necessary to cover themajority (90%) of the population decreases.

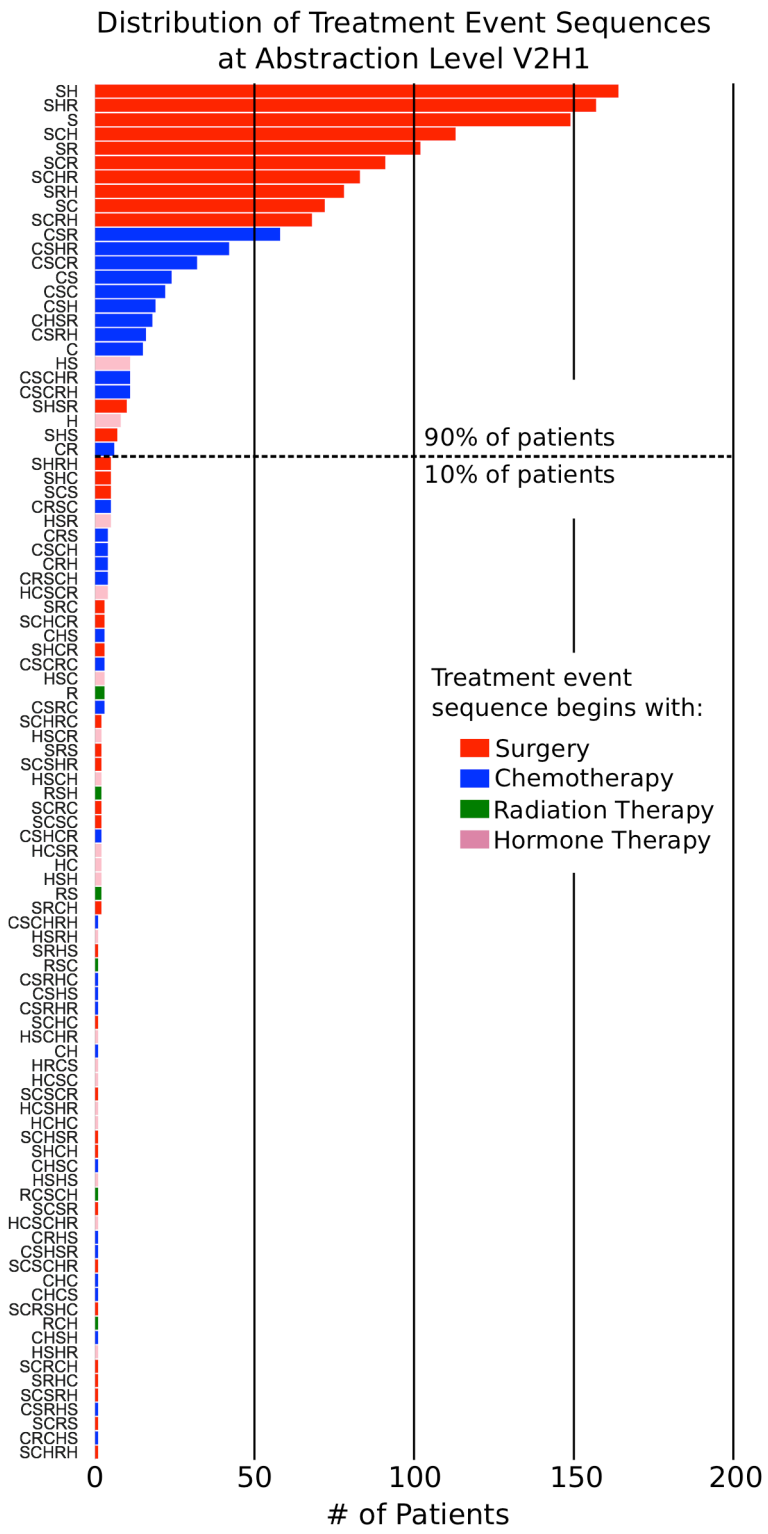

Figure 12: Distribution of treatment event sequences across patients at abstraction level V2H1

The number of patients per treatment event sequence demonstrates an exponential distribution . The first 10 most frequent sequences begin with a surgical treatment event.

To assess the complexity of the treatment event sequences and extent of the horizontal abstraction process, we calculated the number of treatment events per sequence over time for stage I-III patients (Figure 13). The level V0H0 event sequences consistently had almost one additional treatment event in comparison to the level V2H2 treatment string representation. There is no major difference in the average number of treatment events per plan over time.

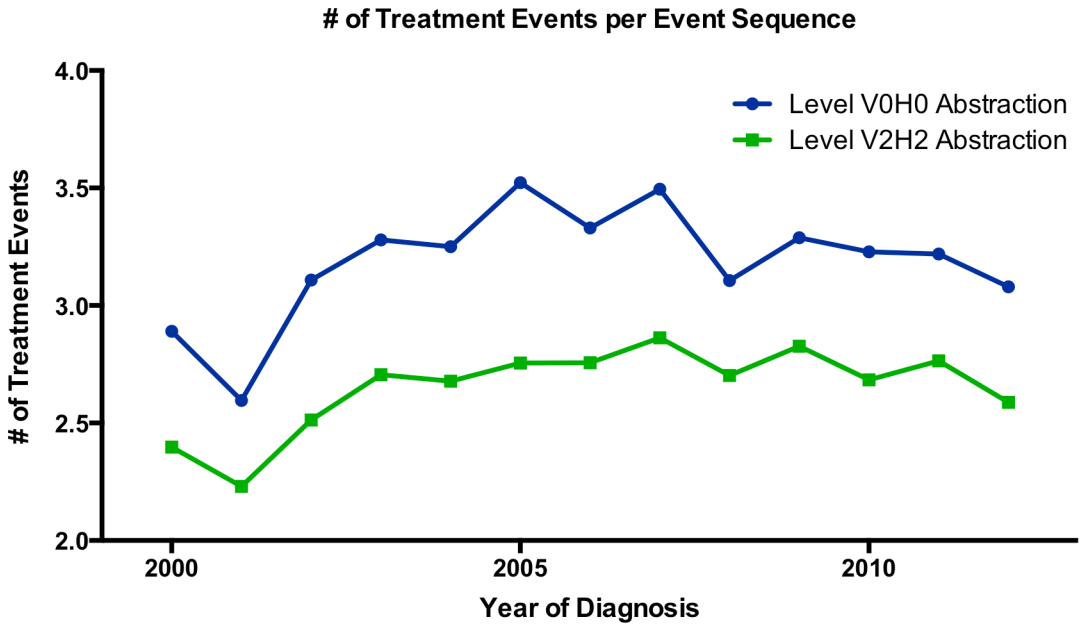

Figure 13: Number of treatment events per event sequence for highest and lowest levels of abstraction

The number of treatment events per treatment plan represents the impact of the horizontal abstraction process. The average number of treatment events per treatment event sequence from 2000-2012 has stayed relatively constant.

We also assessed how the horizontal abstraction process affected the treatment event sequence length based on the number of events in the unabstracted event sequence. We calculated the difference in number of events for the V2H1 and V2H2 levels of horizontal abstraction in comparison to the V2H0 level. The mean and standard deviations at each original event sequence length is presented in Figure 14. The V2H1 level of abstraction appears to

contribute significantly to the consolidation of event sequences in the horizontal abstraction process.

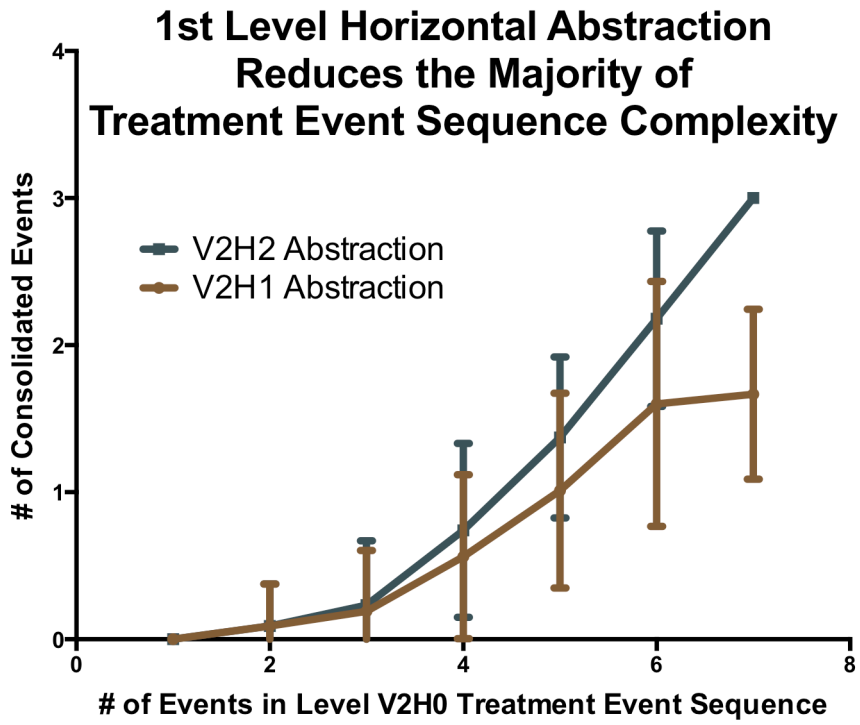

Figure 14: Number of consolidated events in Horizontal Abstraction subtask

Distribution of number of events that were consolidated through the horizontal abstraction process based on length of the original V2H0 level event sequence. Error bars denote standard deviations.

We also developed Sankey diagrams for stage I-III breast cancer patients using our three horizontal abstraction levels to demonstrate visually how treatment plans for a patient population can be organized more compactly at a more abstracted level (Figure 15). Of the 1528 stage I-III patients, 1147 began their course of care with surgery, 321 with chemotherapy, 50 with hormone therapy, and 10 with radiation therapy. The bottom of Figure 15 represents the V1H0 abstraction level that shows the full event sequence with separate mastectomies and breast conserving surgeries; the middle of Figure 15 represents the V2H0 abstraction level that

shows the full event sequences while consolidating mastectomies and breast conserving surgeries; the top of Figure 15 represents the V2H2 abstraction level that only shows the first occurrence of each treatment event type. The transition from top to bottom of Figure 15 demonstrates the consolidation of treatment event sequences to simplify the treatment plan representation while providing a different clinical meaning. We also counted the number of nodes and linkages in the Sankey diagrams when including all treatment event positions in event sequence (Table 3). There appears to be a linear relationship across the various levels of vertical and horizontal abstraction (Figure 16). This indicates that for every new node in a Sankey diagram, there are a stable number of added linkages across abstraction levels. The abstraction process reduces the number of nodes by a factor of 6 and linkages by a factor of 8. The largest decrease in number of nodes and linkages comes in the first horizontal and first vertical abstraction subtask.

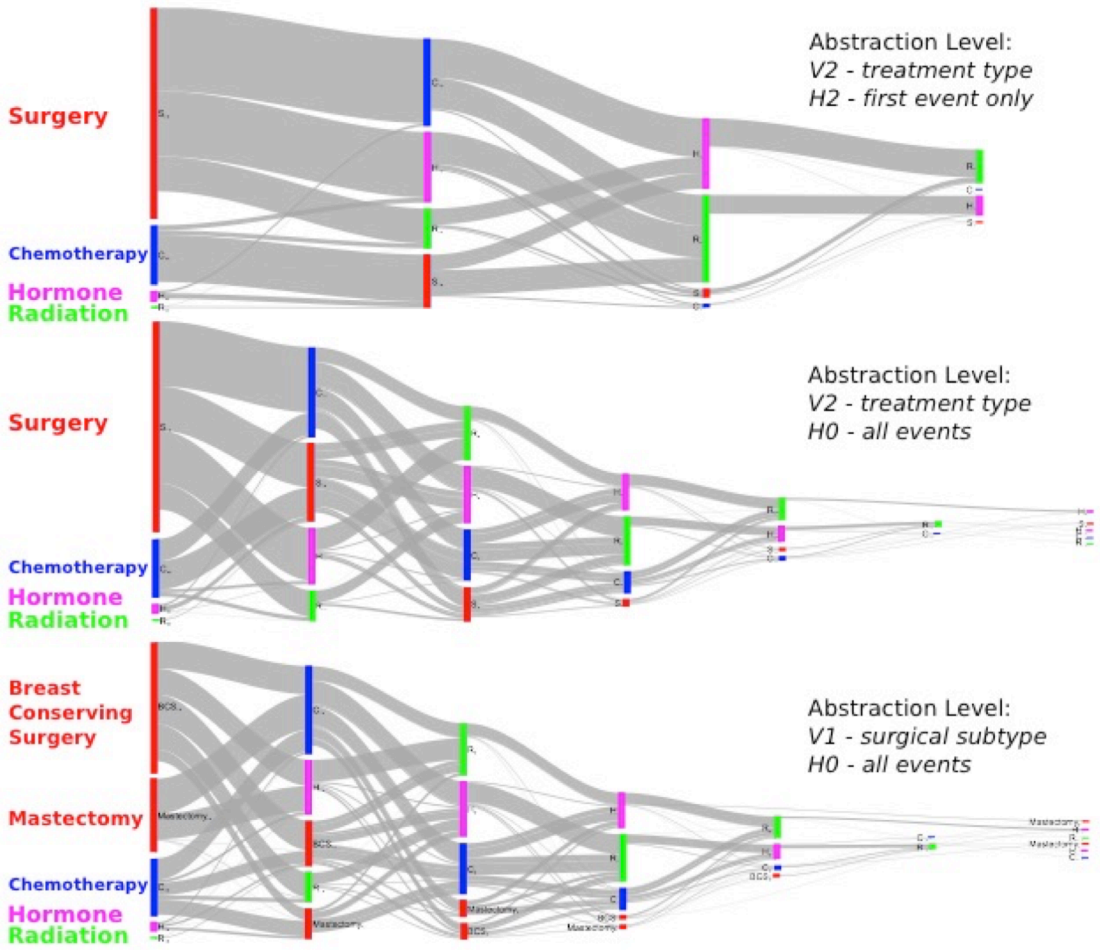

Figure 15: Sankey diagrams for stage I-III breast cancer patients' event sequences at the V1H0, V2H0, & V2H2 levels of abstraction

Sankey diagrams visually representing the distribution of treatment event sequences of 1528 stage I-III breast cancer patients at three abstraction levels. Each vertical line represents a subsequent treatment event with the left-most line representing the patient population. The thickness of the grey bars represents the fraction of the population in that section of the treatment event path.

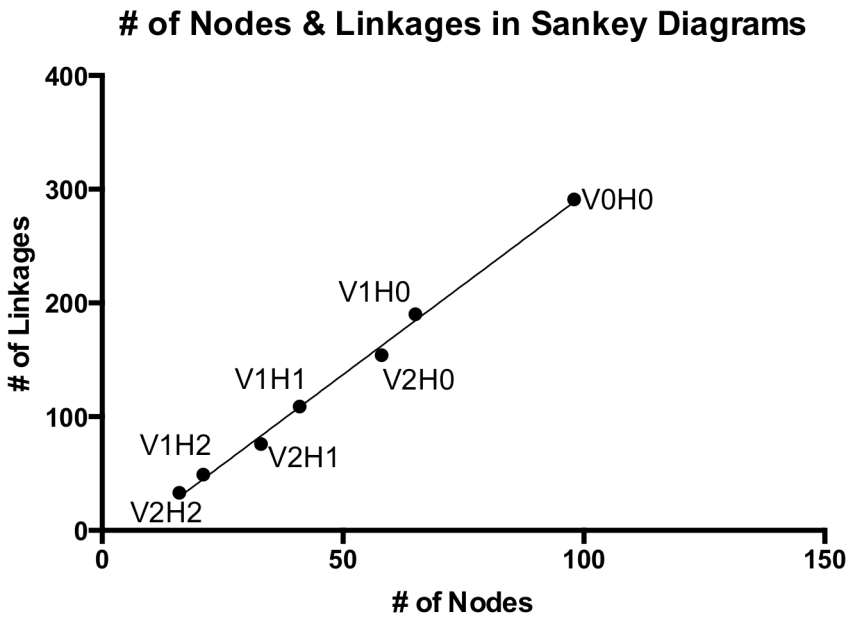

Figure 16: Relationship between number of nodes and links in Sankey diagrams at various abstraction levels

The number of nodes and linkages in the Sankey diagrams at various levels of abstraction demonstrate a linear relationship.

Table 3: Number of nodes and links in Sankey diagrams at various abstraction levels

Abstraction Level	Number of Nodes	Number of Linkages
Horizontal Abstraction		
V1H0	65	190
V1H1	41	109
V1H2	21	49
V2H0	58	154
V2H1	33	76
V2H2	16	33
Vertical Abstraction		
V0H0	98	291
V1H0	65	190
V2H0	58	154

Number of nodes and linkages in Sankey diagrams representing the treatment event sequences of 1528 stage I-III breast cancer patients at various levels of abstraction.

Our final step in characterizing the treatment event sequences was to use the CSPADE event sequence mining method on our treatment event sequences for stage I-III breast cancer patients at the V0H0, V1H0, V1H1, V1H2, V2H0, V2H1, and V2H2 levels of abstraction. This

method yielded 355 frequent event patterns of at least two events at the V0H0 level and 45 patterns at the V2H2 level. When limited to event patterns of at least three events, the V0H0 event sequences resulted with 305 patterns and V2H2 had 30. The list of top frequent event patterns at the V1H0 abstraction level are presented in Table 4 and do include overlap across the patterns. The most frequent event pattern of at least two events is breast conserving surgery followed by a radiation therapy. This shows how event sequence mining could be used to generate the clinical patterns that serve as an input of quality metrics in our Pathfinder method. Additionally, a number of adjuvant and neo-adjuvant event patterns present in the list and show how this can be another valuable method for consolidating patient subgroups. In Figure 17, we show the distribution of treatment event sequences and frequent event patterns at the V1H0 abstraction level. The top 32 event patterns of at least two events cover 90% of the patient population in comparison to the top 126 event sequences. This shows how the event sequence mining method could be used to identify larger patient subgroups. After sorting the event patterns by their support value (frequency), we determined the cumulative fraction of treatment event sequences and patients (Figure 18) that each additional event pattern was present in. The top of Figure 16 demonstrates that the top 50 event patterns of at least two events cover 90% to 97% of all treatment event sequences at all levels of abstraction. The bottom of Figure 16 shows that patient coverage reaches approximately 90% for patterns of at least 2 events.

Table 4: Frequent event patterns at abstraction level V1H0

Frequent Event Pattern	# of Patients	% of Patients
H	918	60.1%
BCS	895	58.6%
R	881	57.7%
C	816	53.4%
Mastectomy	695	45.5%
BCS,R	681	44.6%
C,R	502	32.9%
BCS,H	501	32.8%
C,H	440	28.8%
Mastectomy,H	395	25.9%

H,R	365	23.9%
Mastectomy,C	318	20.8%
BCS,C	317	20.7%
BCS,H,R	266	17.4%
BCS,C,R	247	16.2%
R,H	206	13.5%
C,Mastectomy	188	12.3%
Mastectomy,R	183	12.0%
Mastectomy,C,H	178	11.6%
C,H,R	174	11.4%
BCS,C,H	166	10.9%
BCS,R,H	164	10.7%
C,BCS	142	9.3%
C,BCS,R	119	7.8%
C,R,H	114	7.5%
C,Mastectomy,R	101	6.6%
Mastectomy,C,R	97	6.3%
BCS,C,H,R	78	5.1%
Mastectomy,H,R	76	5.0%
C,Mastectomy,H	71	4.6%
BCS,C,R,H	63	4.1%
C,BCS,H	60	3.9%
H,Mastectomy	45	2.9%
H,BCS	45	2.9%
H,C	44	2.9%
C,BCS,H,R	34	2.2%
H,BCS,R	34	2.2%
Mastectomy,C,H,R	33	2.2%
C,Mastectomy,H,R	33	2.2%
R,C	33	2.2%
Mastectomy,R,H	30	2.0%
Mastectomy,C,R,H	24	1.6%
C,BCS,R,H	19	1.2%
C,H,Mastectomy	19	1.2%
H,C,R	19	1.2%
R,BCS	17	1.1%
H,Mastectomy,R	14	0.9%
C,Mastectomy,R,H	13	0.9%
R,Mastectomy	13	0.9%
BCS,H,C	13	0.9%
C,H,BCS	12	0.8%
C,H,BCS,R	11	0.7%
C,R,Mastectomy	11	0.7%
BCS,R,C	11	0.7%
Mastectomy,H,C	10	0.7%
R,C,H	10	0.7%
C,H,Mastectomy,R	9	0.6%
H,Mastectomy,C	9	0.6%
C,R,BCS	8	0.5%
H,C,BCS	7	0.5%
BCS,H,C,R	7	0.5%
H,C,BCS,R	6	0.4%
R,BCS,H	5	0.3%
H,C,Mastectomy	4	0.3%
R,Mastectomy,H	3	0.2%
C,R,Mastectomy,H	3	0.2%

The number and fraction per frequent event patterns at the level V1H0 abstraction level ranked by frequency.

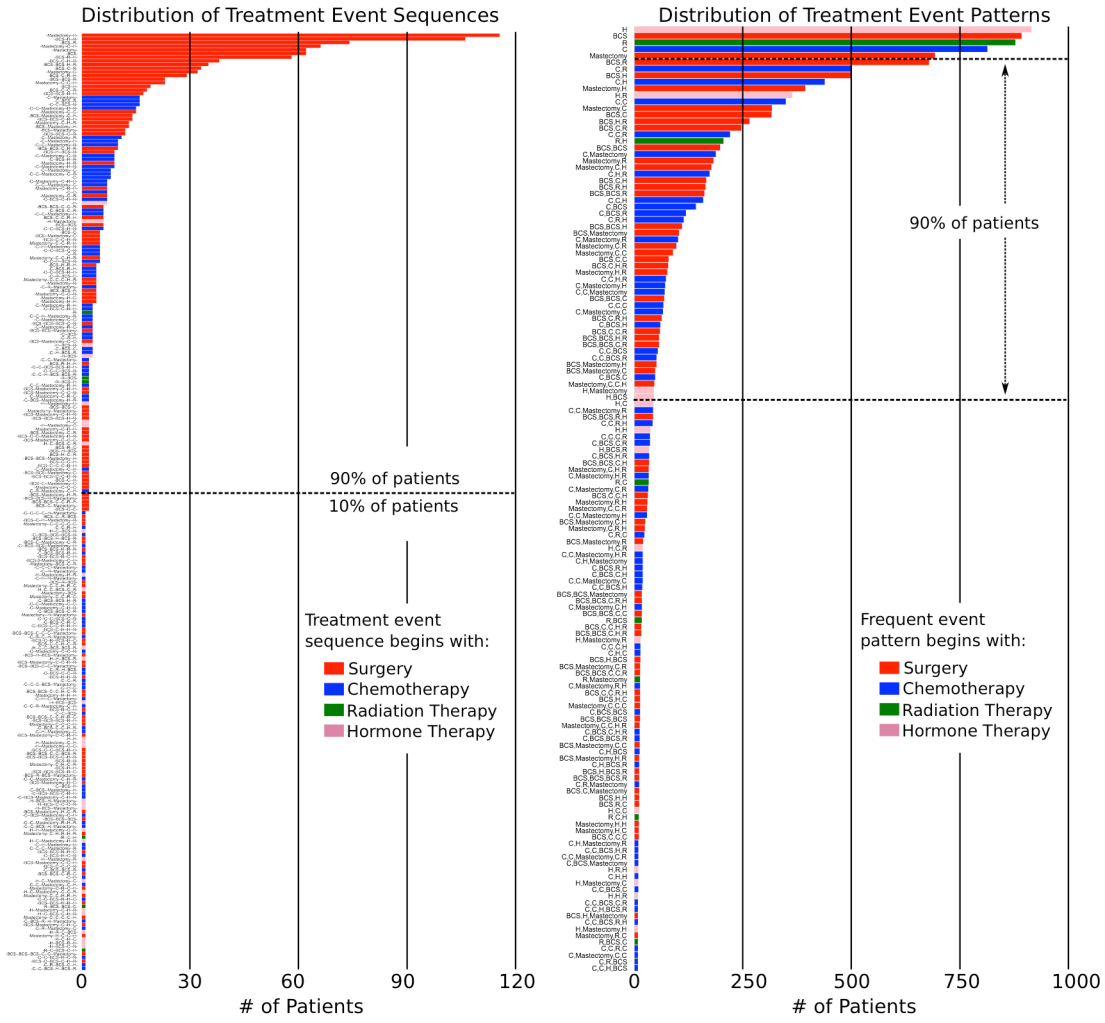

Figure 17: Distribution of treatment event sequences (left) and frequent event patterns (right) at abstraction level V1H0

Left: The number of patients per treatment event sequence demonstrates an exponential distribution . Of the 265 unique treatment event sequences, the top 126 are able to cover 90% of the patient population.

Right: The number of patients per frequent event pattern demonstrates that the 144 most frequent event patterns (of at least 2 events) can cover 90% of the patient population.

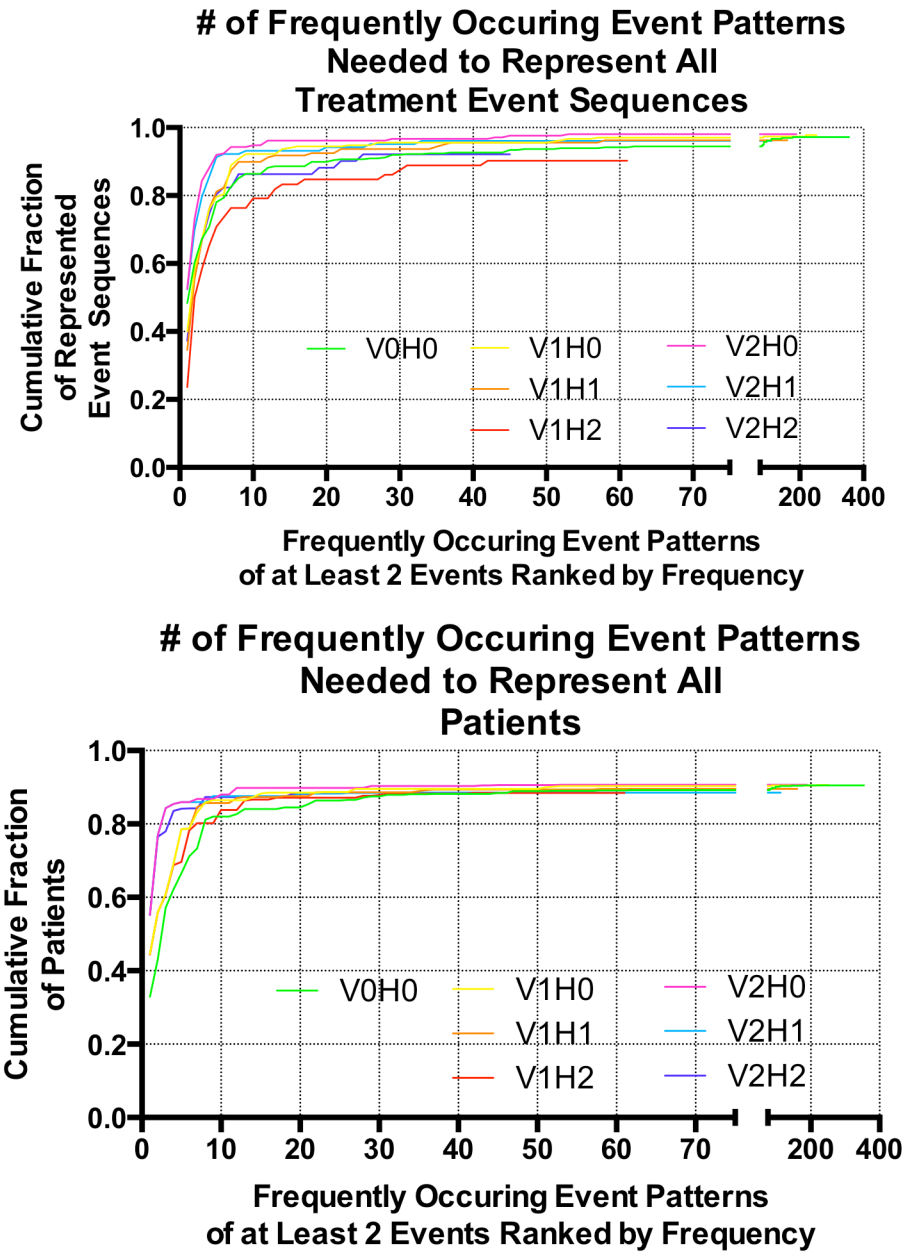

Figure 18: Coverage of frequently occurring event pattern across all treatment event sequences (top) and patients (bottom) at various abstraction levels

The cumulative fraction of treatment event sequences (top) and patients (bottom) represented by each subsequent frequently occurring event pattern at various levels of abstraction. Frequently occurring event patterns are ranked by their frequency and the top represents patterns of at least two events.

5.3: Breast Cancer Treatment Quality Metrics from the Cancer Registry

We next conducted our quality metric querying and generation/visualization subtasks to assess our re-excision event, radiation after breast conserving surgery, and chemotherapy usage quality metrics.

Our first task was to assess the rate of the type of a patients' first surgery at the V1H0 level of abstraction to elucidate the presence of a breast conserving surgery or mastectomy. Figure 19 illustrates the rate of patients initially having either a breast conserving surgery or mastectomy over time. There were 892 patients that received a breast conserving surgery and 592 patients that underwent a mastectomy. The error bars represent the 95% confidence intervals. This analysis demonstrates how rates for mastectomy are increasing compared to breast conserving surgery for a patients' initial surgery.

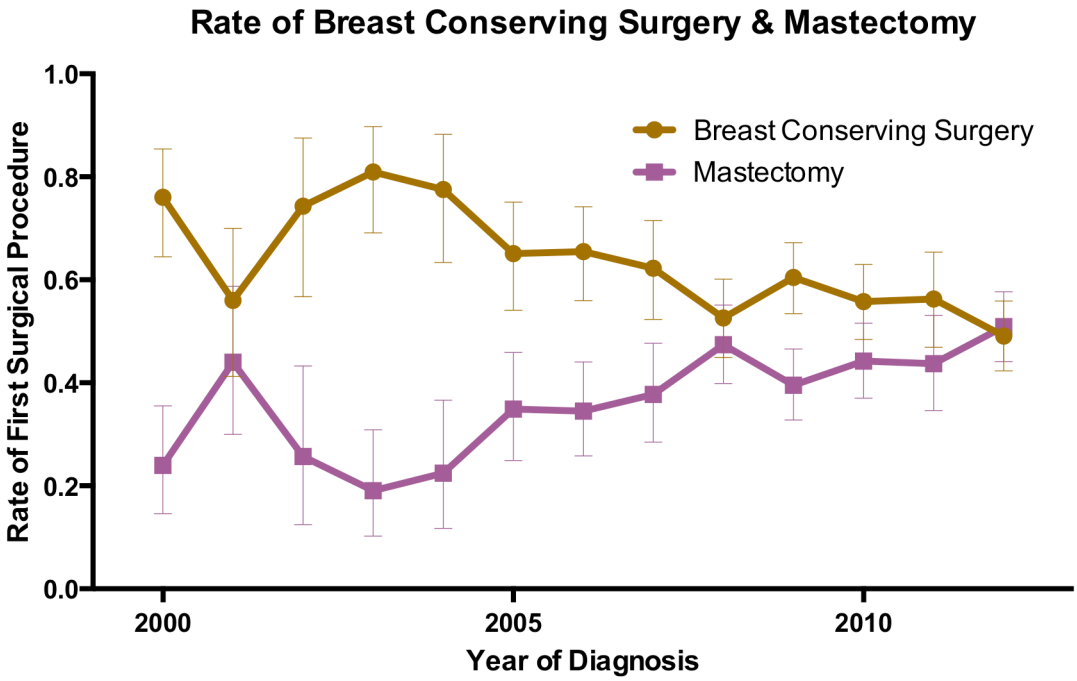

Figure 19: Rate of breast conserving surgery and mastectomy as the first surgical treatment in 1484 stage I-III breast cancer patients

Comparison of patients that begin their surgical course with mastectomy or a breast conserving surgery with error bars denoting the 95% confidence interval. Rates of mastectomy have been increasing over time.

We next assessed the quality of care for patients that began their surgical course with a breast conserving surgery. From this group, Figure 20 represents the rates of patients who required a re-exicision event, progressed to a full mastectomy, or did not require any re-excision at all. Of 890 patients, 608 patients (68%) did not require any following surgery, 181 patients (20%) required a re-excision, and 101 patients (11%) advanced to receive a mastectomy. The error bars represent the 95% confidence intervals. Over time, the need for re-excision or escalation to a mastectomy has decreased with approximately 80% of patients requiring no re-excision or mastectomy. In addition to assessing the rates of additional surgical events after an initial breast conserving surgery, we developed Sankey diagrams at the abstraction level 1V0H (Figure 21). This Sankey diagram illustrates the progression of the surgical course for patients who have a breast conserving surgery.

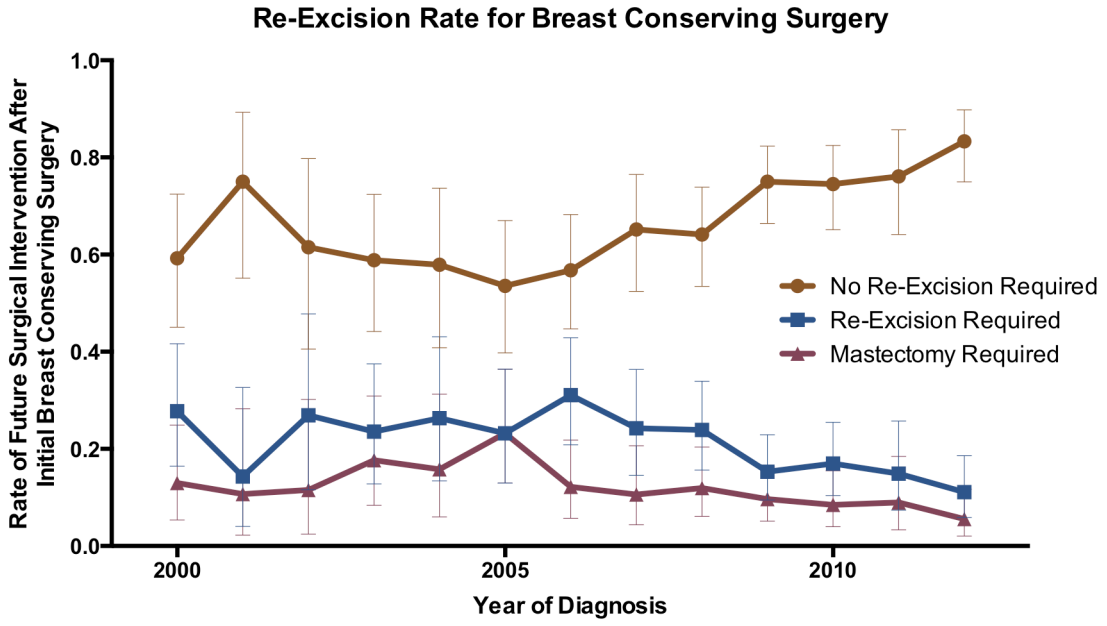

Figure 20: Re-Excision rate for stage I-III breast cancer patients after an initial breast conserving surgery

Rate of re-excision for patients with stage I-III breast cancer that received a re-excision event as their first surgery. The re-excision rate over time is statistically significantly different.

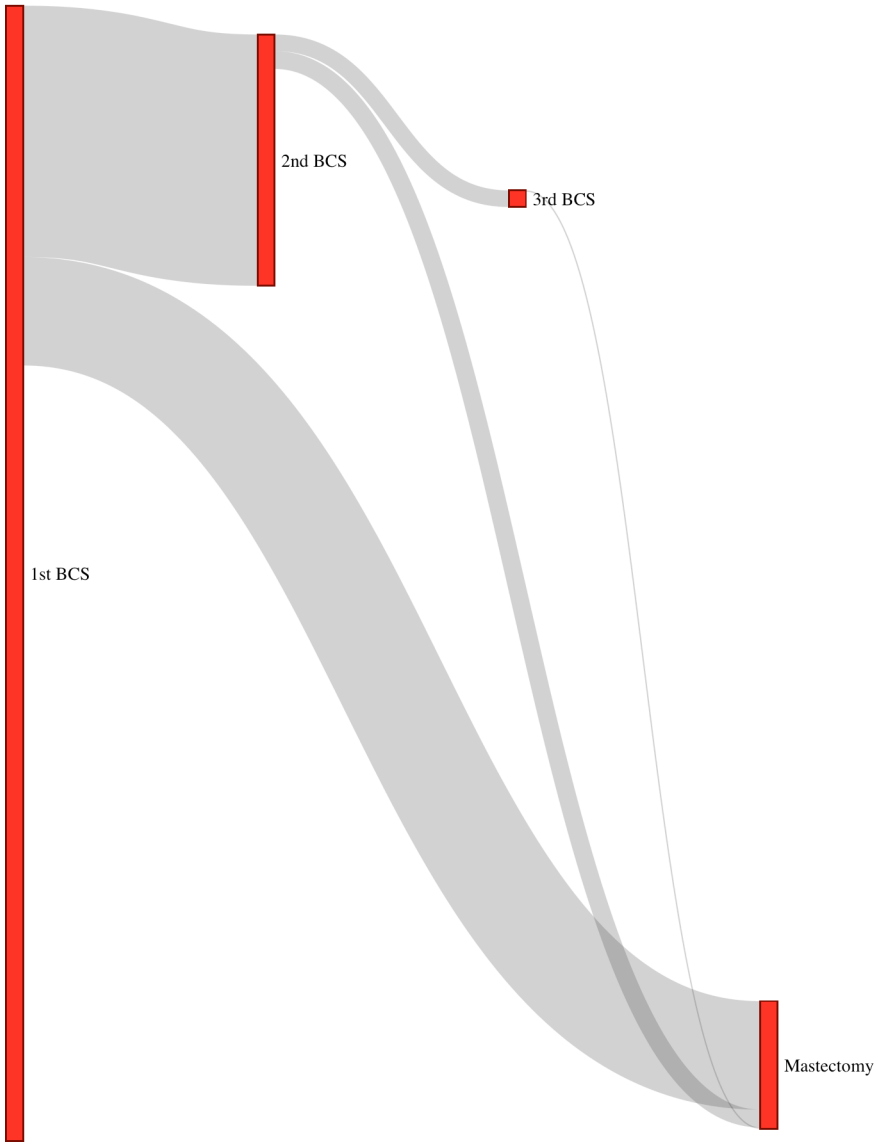

Figure 21: Sankey diagram of re-excision rate after breast conserving surgery

This Sankey diagram visually illustrates the re-excision rate quality metric by demonstrating the need for additional surgical intervention following a breast conserving surgery in 890 stage I-III breast cancer patients. The 890 patients are represented by the height of the left-most vertical blue line. The thickness of the grey bars represents the fraction of those patients that undergo the subsequent segments of the treatment event sequence.

We next analyzed the rate at which patients received radiation therapy after undergoing breast-conserving surgery (Figure 22). The rate appears to have increased from nearly 60% in 2000 to approximately 90% in the mid to late 2000s. In total, 666 patients had radiation therapy after their BCS while 123 did not. The error bars represent the 95% confidence intervals. The rate of radiation therapy after breast conserving surgery has increased and stabilized between 2000 and 2012.

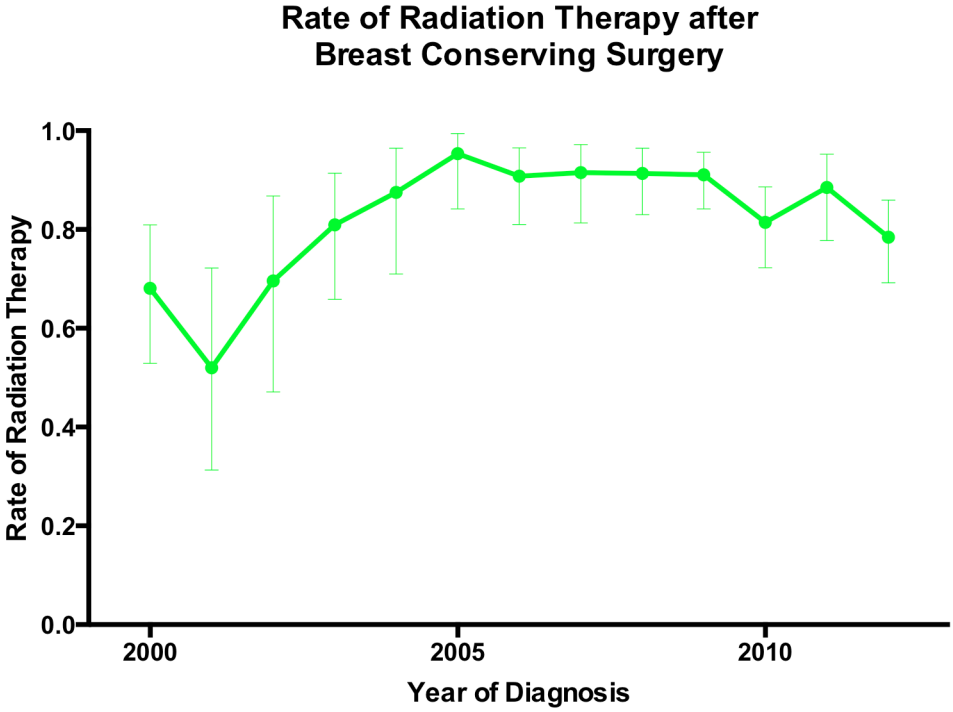

Figure 22: Rate of radiation therapy after breast conserving surgery

Rate at which patients undergoing a breast conserving surgery also receive radiation therapy. The rate of patients receiving radiation therapy over time is statistically significant.

Finally, we assessed the usage of chemotherapy. First we analyzed the rates of neo-adjuvant and adjuvant chemotherapy for stage II and stage III breast cancer patients who are presumably hormone receptor negative as they were not on hormone therapy. Figure 23 indicates the rates at which patients are receiving no chemotherapy, adjuvant chemotherapy, neo-adjuvant chemotherapy, or only chemotherapy. Out of 278 patients, 135 patients (49%) underwent neo-adjuvant chemotherapy, 90 patients (32%) had adjuvant chemotherapy, 32 patients (11%) did not undergo chemotherapy, and 21 patients (8%) had only chemotherapy. The error bars generated via Fisher's exact test represent the 95% confidence intervals. The Chi-square test did not show any statistically significant differences between any combination of the rates of adjuvant, neo-adjuvant, and no chemotherapy over time.

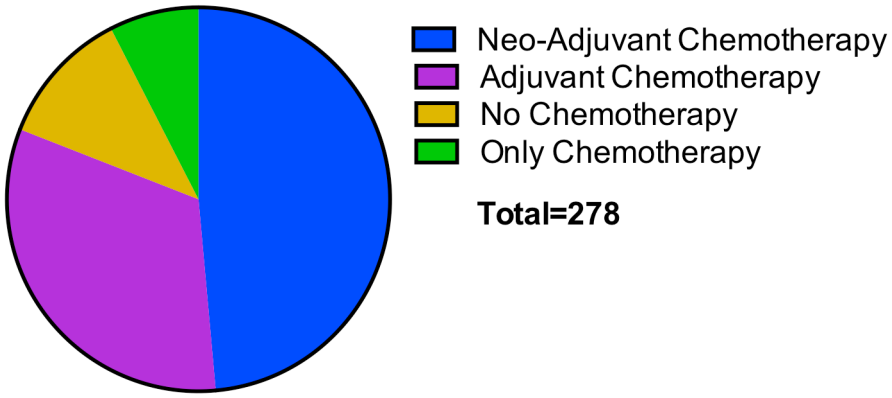

Figure 23: Rate of adjuvant and neo-adjuvant chemotherapy

Fraction of patients that undergo neo-adjuvant chemotherapy (surgery after chemotherapy) in comparison to adjuvant chemotherapy (surgery before chemotherapy) for Stage II and Stage III breast cancer. The rates were not statistically significantly different over time.

We finally looked at hormone receptor positive, HER2 receptor negative, lymph node negative, stage I-II breast cancer patients to determine if the OncotypeDX test had altered the usage of chemotherapy (Figure 24). Out of 695 patients, 419 patients (60%) did not receive chemotherapy while 276 (40%) did. The error bars represent the 95% confidence intervals. Chemotherapy usage seems to have declined only slightly since 2005.

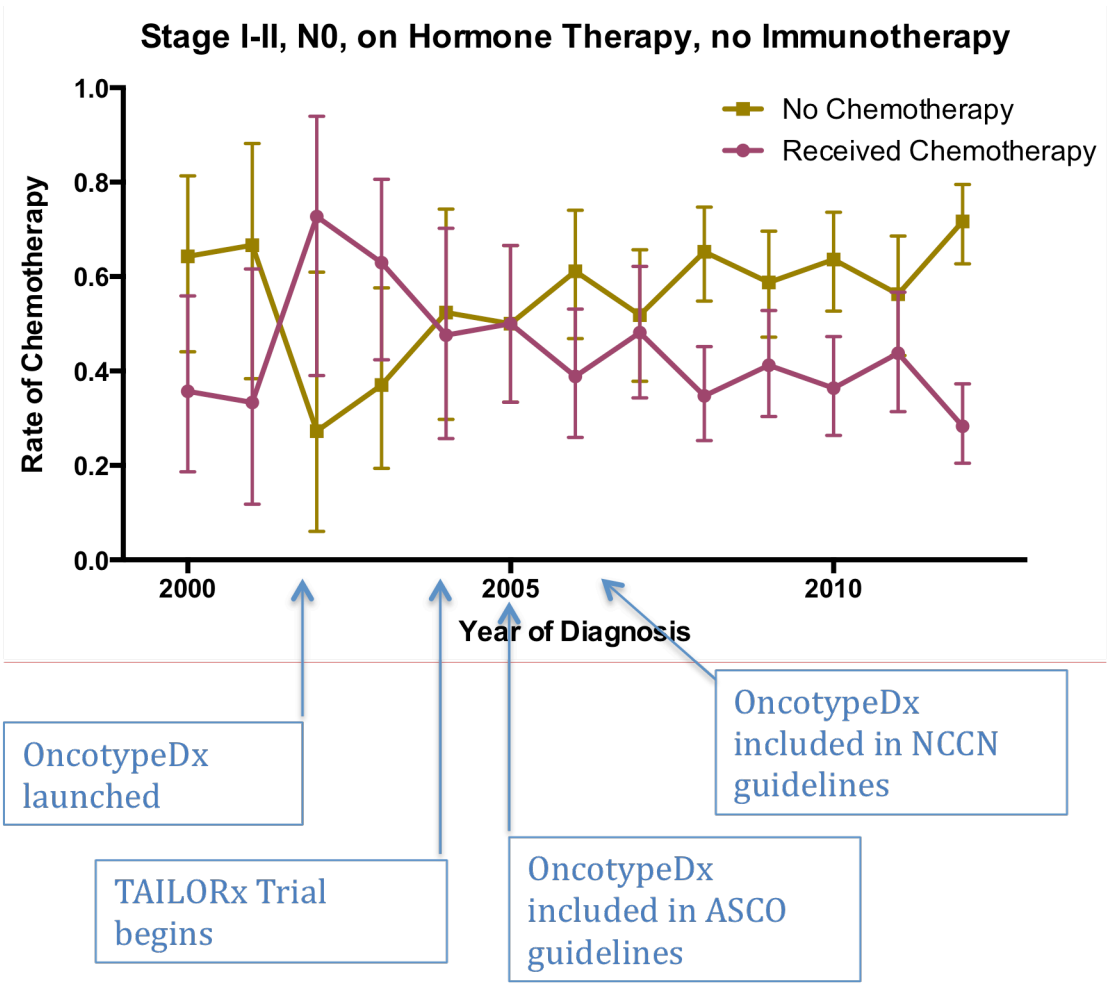

Figure 24: Rate of patients receiving chemotherapy due to the OncotypeDx test

Chemotherapy usage for hormone receptor positive, HER2 receptor negative, lymph node negative, stage I-II breast cancer patients. A timeline of OncotypeDx usage is included to demonstrate how OncotypeDx was launched in 2002 and incorporated into guidelines in 2005-2006. The rate of chemotherapy over time was statistically significant.

5.4: Breast Cancer Treatment Plan Abstraction from Surgical CPT Codes

Our next objective was to assess the viability of using surgical CPT codes in the abstraction and quality metric process. Breast surgical CPT codes were extracted that occurred within one day of a cancer registry surgical event. We defined breast surgical CPT codes as CPT codes between 19000 and 19999; they represent events such as mastectomies and breast conserving surgeries as well as other breast events (i.e., aspirations, cyst excisions, breast localization device placements). Of the 2679 patients in our cohort, 2524 patients had 3208 surgical events in the cancer registry. We first assessed the number of cancer registry surgical events that did not have an associated breast surgical CPT code in the VUMC clinical information system (Figure 25, Table 5). Of the 2420 surgeries that took place at VUMC, 97.9% had a breast surgical CPT code and 92.5% had a breast conserving surgery or mastectomy CPT code occurring within one day of the surgical event in the VUMC system. Of the 788 surgeries that took place elsewhere, only 13.7% had a breast surgical CPT code present and only 10.5% had a specific mastectomy or breast conserving surgery CPT code; many of the CPT codes present were from Williamson County, a local community hospital affiliate where some VUMC providers practice.

Table 5: Number of missing CPT codes for surgical events in the cancer registry

Breast Surgical CPT Code

	At Vanderbilt	Not at Vanderbilt	Total
CPT Present	2369	108	**2477**
CPT Not Present	51	680	**731**
Total	**2420**	**788**	**3208**

Mastectomy or Breast Conserving Surgery CPT Code

	At Vanderbilt	Not at Vanderbilt	Total
CPT Present	2238	83	**2321**
CPT Not Present	182	705	**887**
Total	**2420**	**788**	**3208**

Number of cancer registry surgical events conducted at or outside of Vanderbilt where a CPT code was or was not present within one day of the surgical event. The top table focuses on all breast cancer surgical CPT codes while the bottom table looks at specific breast conserving surgery and mastectomy CPT codes.

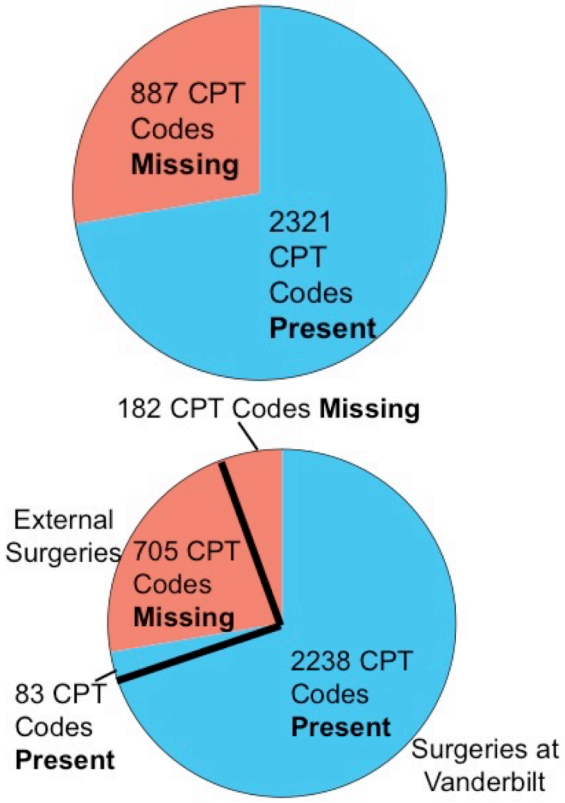

Figure 25: Fraction of missing CPT codes for surgical events in the cancer registry

Top: Fraction of cancer registry surgical events that had corresponding CPT codes present in the administrative data. The corresponding CPT code for a mastectomy or breast conserving surgery was required to take place within 1 day of the registry event date.
Bottom: Same figure as on top but with the 2420 surgeries that took place at Vanderbilt and 788 surgeries that were conducted externally.

Of the 887 missing mastectomy and breast conserving surgery CPT codes, the largest fraction of missing codes were for the first surgery in the event sequence 87% (771), followed by 12% (108) in the second and 1% (8) in the third (Table 6, Figure 26). We then analyzed missing surgical CPT codes on a per person basis by measuring the number of patients that had a specific mastectomy or breast conserving surgery CPT code present for every surgical treatment event (Figure 27). Of the 2524 patients, 793 patients (31.4%) were missing at least one CPT code while 1731 patients (68.6%) did not have any missing codes. Finally, we measured how well the CPT codes aligned with the cancer registry codes (Table 7) using a per-event and per-patient analysis. For the per-event analysis, we calculated the alignment of the 2321 cancer registry surgical events with their respective breast conserving surgery or mastectomy CPT codes at the V0, V1, and V2 levels of abstraction. When represented at the V0 level of abstraction, 57% of registry codes aligned with CPT codes at their granular level. The level of alignment increased to 97% at the V1 and 99% at the V2 levels of abstraction. A similar trend exists with per-patient alignment indicating that the codes are well aligned when comparing breast conserving surgery and mastectomy.

Table 6: Fraction of missing CPT codes per surgical event position in the cancer registry

Position in Surgical Sequence	Total Number of CPT Codes	Number of Missing CPT Codes	Fraction of Missing Codes in Position
1	2830	771	27.2%
2	612	108	17.6%
3	135	8	5.9%
4	17	0	0.0%
5	1	0	0.0%

Number of missing specific breast conserving surgery and mastectomy CPT codes at different positions in the surgical event sequence. A CPT code was present if it occurred within one day of a cancer registry surgical event. The highest fraction of missing codes is for the first surgery of a surgical sequence.

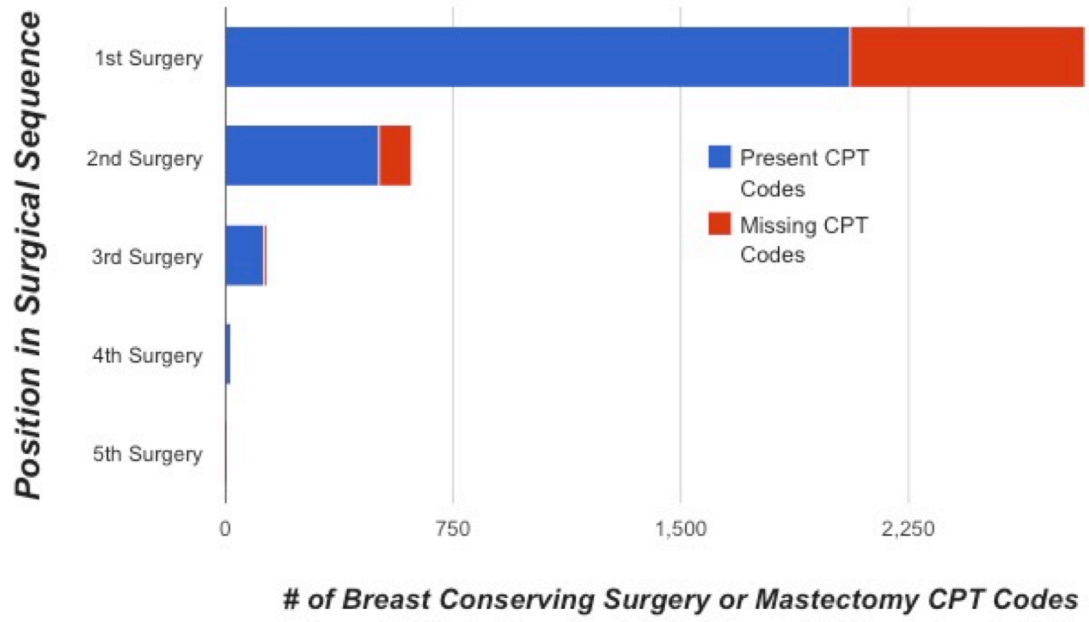

Figure 26: Fraction of missing surgical CPT codes per surgical event position in cancer registry

The majority of missing CPT codes are from the first surgical event in the event sequence.

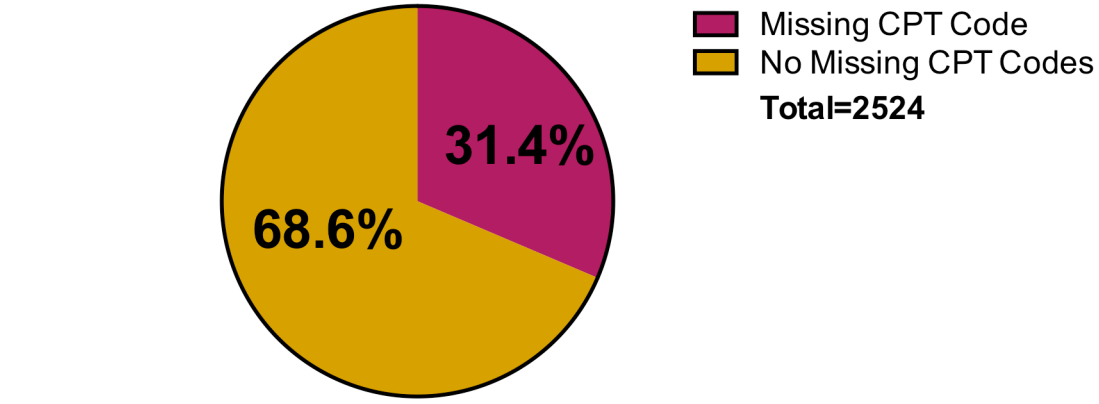

Figure 27: Fraction of breast cancer patients with at least one missing surgical CPT code

Fraction of patients that have a CPT code within one day of a cancer registry surgical event for each event in a patient's surgical event sequence.

Table 7: Fraction of surgical CPT codes aligned the cancer registry surgical codes

Per Event

	V0	V1	V2
Aligned	1322 (57.0%)	2257 (97.2%)	2320 (99.96%)
Misaligned	999 (43.0%)	64 (2.8%)	1 (0.04%)

Per Person

	V0	V1	V2
Aligned	1063 (53.0%)	1941 (96.8%)	2004 (99.95%)
Misaligned	942 (47.0%)	64 (3.2%)	1 (0.05%)

Number of cancer registry surgical events that aligned with their respective matched CPT codes. The top table counts the number of surgical events that were misaligned while the bottom table counts the number of patients with at least one misaligned CPT code.

5.5: Breast Cancer Treatment Quality Metrics from Surgical CPT Codes

We finally used the surgical CPT codes and their NCI Thesaurus mapping to reassess two surgical quality metrics. Of the 1528 stage I-III breast cancer patients with a surgical event in our patient cohort, 1282 had at least one CPT code for a breast conserving surgery or mastectomy. We selected CPT codes within 2 years of the original breast surgery code for each patient in the cohort. The year of diagnosis was chosen as the year of the first surgery in the sequence. We first analyzed the rate of breast conserving surgery and mastectomy and overlaid them against the rates generated from the cancer registry data (Figure 28). The error bars represent the 95% confidence intervals. This indicates, like with the cancer registry data, that not only are the rates between breast conserving surgery and mastectomy statistically significantly different, but also that there is a significant linear trend over time.

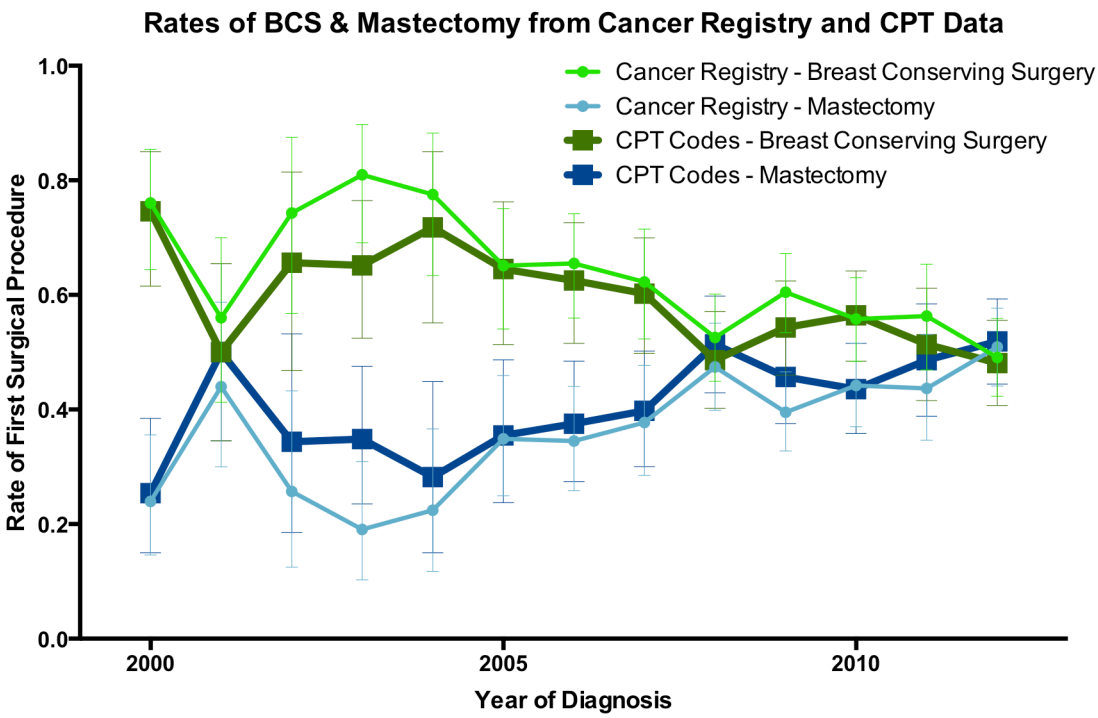

Figure 28: Rate of breast conserving surgery and mastectomy in breast cancer patients from CPT codes and the cancer registry

Comparison of patients that begin their surgical course with mastectomy or a breast conserving surgery using both CPT and cancer registry data. CPT data is bolded. Error bars denote the 95% confidence interval and show no statistically significant difference between rates for CPT and cancer registry data. Rates of mastectomy have been increasing over time.

We also assessed the re-excision rate using CPT data for patients that began their surgical course of care with a breast conserving surgery (Figure 29). The error bars represent the 95% confidence intervals. These findings are similar to the cancer registry data. Over time, the need for re-excision or escalation to a mastectomy has decreased with approximately 75% of patients requiring no re-excision or mastectomy.

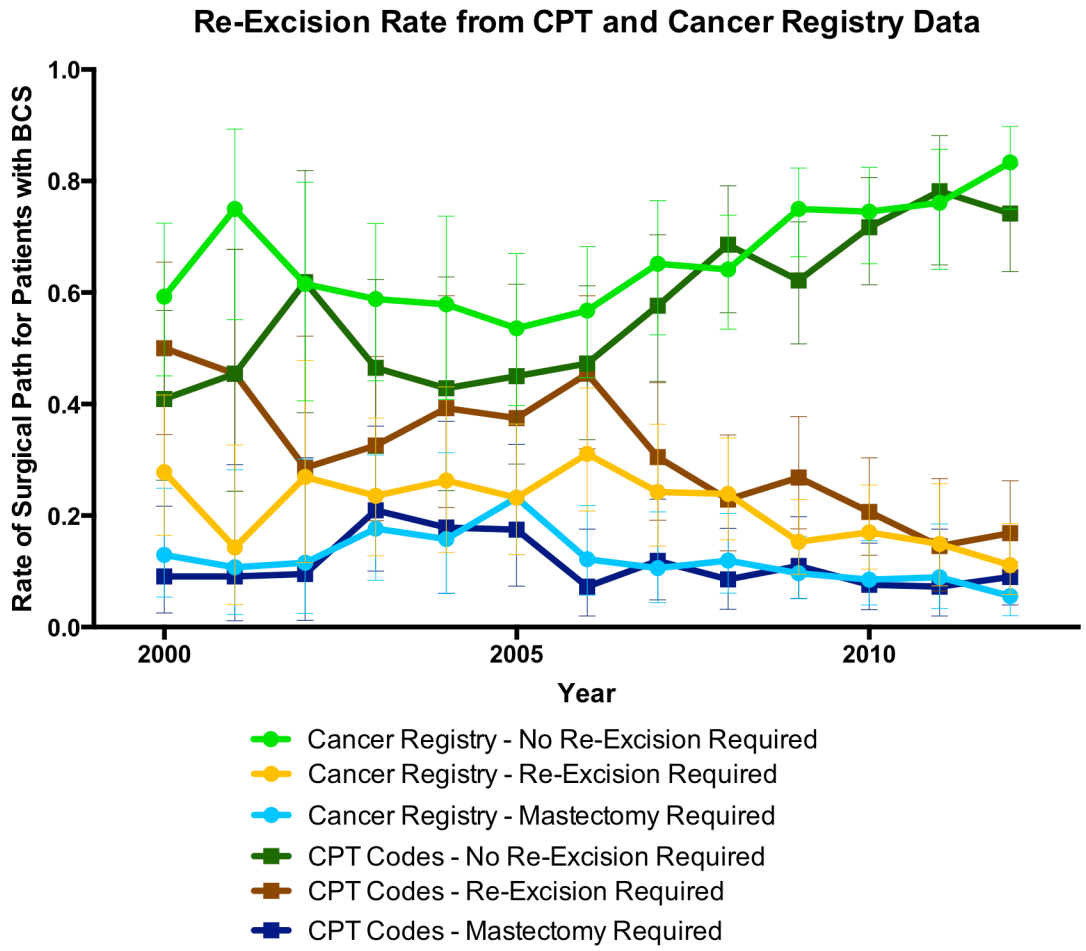

Figure 29: Re-Excision rate in breast cancer patients from CPT codes and the cancer registry

Comparison of patients that begin their surgical course with mastectomy or a breast conserving surgery using both CPT and cancer registry data. CPT data is bolded. Error bars denote the 95% confidence interval and show no statistically significant difference between rates for CPT and cancer registry data. Rate of no re-excision events have been increasing over time.

The alignment of the cancer registry and CPT code abstractions was assessed by comparing surgical event sequences at multiple levels of abstraction (Table 8). The alignment improved at higher levels of abstraction as the amount of treatment event sequence variability was reduced. The largest impact on alignment was on the V0 to V1 and H0 to H1 transformations. The V0 to V1 abstraction primarily involves the conversion of specific surgery

types to either a breast conserving surgery or mastectomy. The H0 to H1 abstraction removes directly adjacent repeating events.

Table 8: Comparison of surgical event sequences from CPT codes and the cancer registry at various abstraction levels

Abstraction Level	Equivalent Abstraction	Different Abstraction
V0H0	587 (45.8%)	695 (54.2%)
V1H0	915 (71.4%)	367 (28.6%)
V2H0	940 (73.3%)	342 (26.7%)
V1H1	1166 (91.0%)	116 (9.0%)
V1H2	1172 (91.4%)	110 (8.6%)

Comparison of surgical event sequences generated from CPT and cancer registry event data at various levels of abstraction.

Finally, we used the event sequence-mining package to identify frequent event patterns from CPT and cancer registry surgical event sequences. We then assessed the number of patients that had each event pattern using both the CPT and cancer registry data (Table 9). We then compared the resulting antecedent and consequent pairs. Three of the five identified event sequence patterns were shown to be statistically significantly different between the two sets despite relatively similar rates.

Table 9: Comparison of the occurrence of frequent surgical event sequences in CPT codes and the cancer registry

	% of Patients with Surgical Event Sequence		
Frequent Event Pattern	**Cancer Registry**	**CPT Data**	**p-value**
BCS - BCS	14.1%	10.9%	0.017*
Mastectomy - Mastectomy	0.2%	2.5%	<0.0001*
BCS - Mastectomy	7.1%	4.8%	0.0151*
BCS – BCS - Mastectomy	1.2%	1.3%	0.9982
BCS – BCS - BCS	0.9%	0.9%	0.830

Fraction of patients with frequent event pattern using both cancer registry and CPT event data. The BCS-BCS, Mastectomy-Mastectomy, and BCS-Mastectomy patterns were shown to be statistically significantly different between the CPT and cancer registry event data.

Chapter 6: Conclusions and Discussion

This study showed how the six subtasks of the Pathfinder temporal abstraction framework could be used to represent treatment event paths and generate clinical quality metrics. We assessed and visualized how our abstraction process could simplify and consolidate our set of treatment event paths. We also demonstrated how administrative data could be used to generate clinical quality metrics in a more near-real time fashion in comparison to cancer registry data. We describe our contributions to the informatics (Section 6.1) and clinical (Section 6.2) domains as well as their corresponding limitations and future directions.

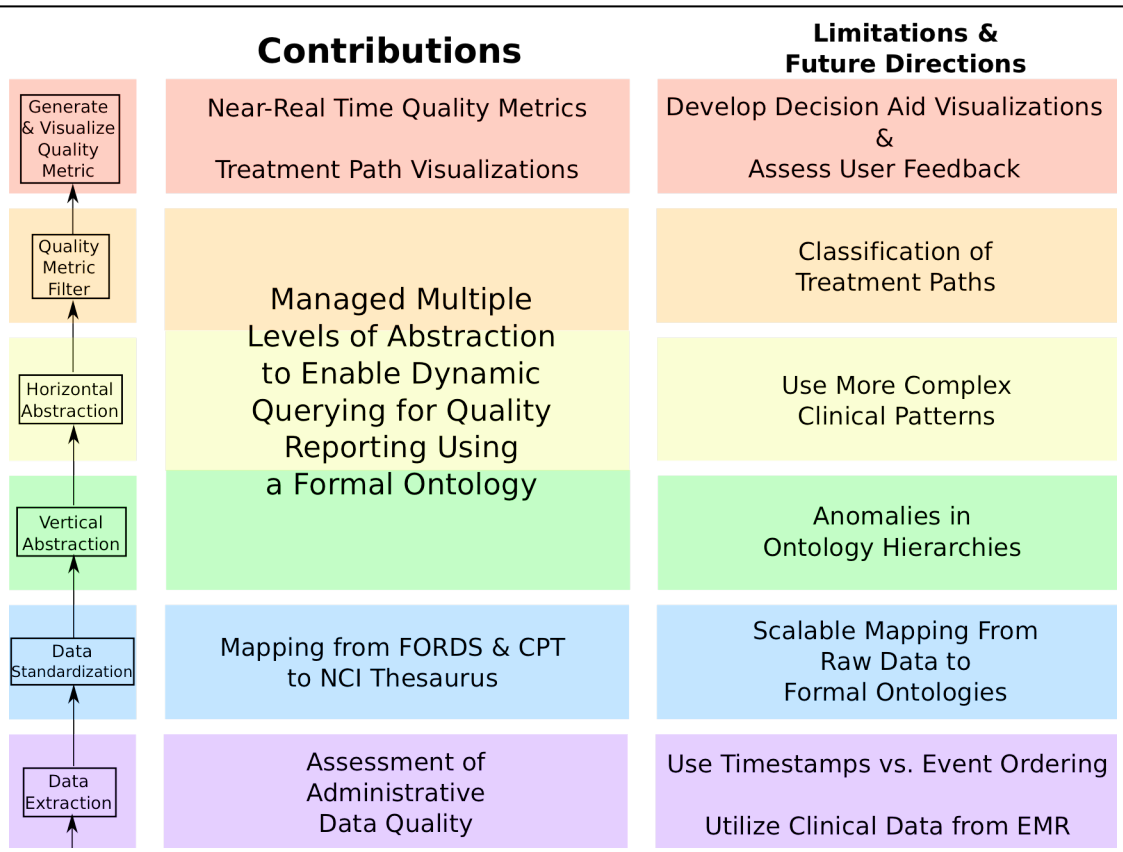

Figure 30: Contributions, limitations, and future directions across our six subtasks

6.1: Informatics Contributions, Limitations, & Future Directions

We have developed a data abstraction framework to map raw treatment events to the level of abstraction necessary to answer a variety of desired quality metrics. We represented the treatment events at multiple levels of abstraction simultaneously and studied the patterns of care across our breast cancer patient population.

In the data extraction and standardization steps, we placed treatment events in order to generate treatment event sequences and mapped the FORDS registry and administrative CPT event codes to their corresponding NCI Thesaurus concepts. We manually developed a mapping between the FORDS and CPT codes for mastectomy and breast conserving surgery because the mappings were not robust in the UMLS and were not present in the NCI Metathesaurus (due to the proprietary nature of FORDS and CPT). Our work was limited because we only used treatment event ordering and we relied on a manually derived mapping. The treatment event sequences were atemporal and did not represent time intervals. This prevented users from querying for event time intervals or understanding length of time in treatment event paths. Furthermore, the manual mapping limited this work to focus specifically on the domain of surgical events that were mapped. In the future we will need to represent time in our framework and develop scalable methods to develop mappings from treatment events to a formal ontology. We could utilize symbolic temporal logic to track time stamps and intervals for various event-based data. We could potentially use natural language processing and existing mapped events to generate any required mappings in a more scalable fashion.

In the vertical and horizontal abstraction steps, we managed treatment event sequences at multiple levels of vertical and horizontal abstraction. We managed the sequences at multiple levels of abstraction simultaneously in order to be able to query them based on the clinical question that was asked. Our work to represent the events at multiple levels of representation

differentiates this work from Shahar's Knowledge Based Temporal Abstraction (KBTA) framework. The KBTA system primarily focused on a direct abstraction path that was focused on a single clinical question. Out system is limited by anomalies in ontology hierarchies and the manual development of clinical patterns. Our vertical abstraction step utilized a subset of the NCI Thesaurus hierarchy that focused on breast conserving surgery. As we expand our treatment event mappings, we will have to deal with the various complexities of ontology hierarchies (i.e., concepts in multiple trees, unnecessary concepts in layers). In the future, we will have to be able to not only develop our treatment mappings in a scalable way but also efficiently extract the hierarchy we need. Finally, we used three simple clinical patterns for our horizontal abstraction step that were relevant to cancer care. In the future, we can incorporate more complex clinical event patterns that we generate from various data mining methods such as process mining.

In the quality metric filtering and generation steps, we dynamically queried our set of abstracted treatment event sequences to generate and visualize various quality metrics. We also showed how administrative data could serve as a data source for generating near-real time quality metrics. This has the potential to allow for clinical quality metrics to be measured on larger patient populations without the labor-intensive manual review process or the six-month delay for the manual curation of cancer registry data. This will allow for quality metrics and other analysis to be delivered to providers, administrators, and patients in near real-time for a larger segment of the patient population. This work is limited by lack of information present in cancer registry and administrative data. Information on complications or adverse events might only be present in the clinical record and will need to be extracted from medical record data. In the future, patients, providers, administrators, and other stakeholders should be able to explore

quality metrics and treatment paths on an interactive, dynamic, and query-able interface which utilizes a temporal abstraction framework.

6.2: Clinical Contributions, Limitations, & Future Directions

We demonstrated that our system could generate a variety of clinical quality metrics from cancer registry data without the need for any additional manual review.

In the data extraction and standardization steps, we assessed the quality of our administrative data in comparison to the cancer registry for breast conserving surgeries and mastectomies. We were able to measure the fraction of missing and misaligned CPT codes. Our efforts are limited by missing data in the Vanderbilt system and by the kind of data stored in the registry and CPT codes. In the future, we need to be able to detect events from the record (i.e., pathology reports, clinical notes) that occurred outside of Vanderbilt. Furthermore, we will need to use additional data sources from the electronic medical record system to reduce the risk of misidentifying events. Another potential data input could include national cancer registry data sets,that while limited in the information they contain would provide a larger nationwide patient population.

In the vertical and horizontal abstraction steps, we abstracted treatment event sequences to many different levels of clinical relevance and significance. We were able to represent treatment paths across the patient population at many levels of abstraction. Our abstraction process was able to consolidate the treatment event sequences of a majority of patients while the long tail of the sequence distribution is not simplified through the abstraction process. We also visually represented the simplification that results from the abstraction process through Sankey diagrams. In the future, additional analysis will be necessary to assess

the relationships between complexity and quality of care, assess outcomes including adverse events, and develop methods to effectively group similar treatment event paths.

In the quality metric filtering and generation steps, we showed that we could generate quality metrics that mirrored national and institutional rates. We also started to develop visualizations of treatment paths and quality metrics. By appropriately querying our abstraction framework, we were able to generate metrics on re-excision rates, radiation after breast conserving surgery, and chemotherapy usage. Our calculated quality metrics showed Vanderbilt matching or surpassing previously reported quality studies. For example, prior studies report a ~22.9% re-excision rate (McCahill et al., 2012) and 25% chemotherapy usage rate for patients who received the OncotypeDX test(Asad et al., 2008). These were similar to Vanderbilt re-excision rate as well as chemotherapy usage rates in patients suspected of recieveing the OncotypeDX test. Additionally, our rate of radiation therapy after breast conserving surgery is slightly under the rate identified by a review of the Vanderbilt cancer registry for patients diagnosed in 2011. We have also demonstrated how we can incorporate data visualization into the temporal abstraction framework. We demonstrated how a quality metric such as the re-excision rate could be visualized as a Sankey diagram to show the progression of patients through their treatment. Given the focus on population health management, providing additional visual context for treatment event sequences can be beneficial in comparison to a single quality metric rate. Sankey diagrams and other visual representations could be utilized to serve as a decision aid for patients and can serve as real-time quality feedback for providers. The integration of an interactive visualization system into a real-time quality metric generation temporal abstraction framework could help improve communication and feedback.

REFERENCES

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. Retrieved from http://dl.acm.org/citation.cfm?id=170072

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*. Retrieved from http://www.sciencedirect.com/science/article/pii/0004370284900080

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

Asad, J., Jacobson, A. F., Estabrook, A., Smith, S. R., Boolbol, S. K., Feldman, S. M., … Tartter, P. I. (2008). Does oncotype DX recurrence score affect the management of patients with early-stage breast cancer? *American Journal of Surgery*, *196*(4), 527–9. http://doi.org/10.1016/j.amjsurg.2008.06.021

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. Retrieved from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Applying+Temporal+Abstraction+in+Medical+Information+Systems#0

Barreto, J. A., O'Malley, W., Kubeil, M., Graham, B., Stephan, H., & Spiccia, L. (2011). Nanomaterials: applications in cancer imaging and therapy. *Advanced Materials (Deerfield Beach, Fla.)*, *23*(12), H18–40. http://doi.org/10.1002/adma.201100140

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

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

Braun, S. (2003). The History of Breast Cancer Advocacy. *The Breast Journal*, *9*(s2), S101–S103. http://doi.org/10.1046/j.1524-4741.9.s2.13.x

C.L., V.-J., Berger, M., & Pao, W. (2014). Types of Molecular Tumor Testing - My Cancer Genome. Retrieved September 27, 2014, from http://www.mycancergenome.org/content/other/molecular-medicine/types-of-molecular-tumor-testing/

Carey, L. a, Dees, E. C., Sawyer, L., Gatti, L., Moore, D. T., Collichio, F., … Perou, C. M. (2007). The triple negative paradox: primary tumor chemosensitivity of breast cancer subtypes. *Clinical Cancer Research : An Official Journal of the American Association for Cancer Research*, *13*(8), 2329–34. http://doi.org/10.1158/1078-0432.CCR-06-1109

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11776740

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

Chapman, S. a, Mulvihill, L., & Herrera, C. (2012). Workload and time management in central cancer registries: baseline data and implication for registry staffing. *Journal of Registry Management*, *39*(4), 178–84. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/23493024

Chassin, M. R., & Loeb, J. M. (2011). The ongoing quality improvement journey: next stop, high reliability. *Health Affairs (Project Hope)*, *30*(4), 559–68. http://doi.org/10.1377/hlthaff.2011.0076

Chassin, M. R., Loeb, J. M., Schmaltz, S. P., & Wachter, R. M. (2010). Accountability measures--using measurement to promote quality improvement. *The New England Journal of*

Medicine, *363*(7), 683–8. http://doi.org/10.1056/NEJMsb1002320

Chen, C., & Ackerly, D. C. (2014). Beyond ACOs and bundled payments: Medicare's shift toward accountability in fee-for-service. *JAMA : The Journal of the American Medical Association*, *311*(7), 673–4. http://doi.org/10.1001/jama.2014.11

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

Chernew, M. E., Baicker, K., & Hsu, J. (2010). The specter of financial armageddon--health care and federal debt in the United States. *The New England Journal of Medicine*, *362*(13), 1166–8. http://doi.org/10.1056/NEJMp1002873

Clinical Trials Search Results - National Cancer Institute. (n.d.). Retrieved September 27, 2014, from http://www.cancer.gov/clinicaltrials/search/results?protocolsearchid=13275905

Collins, F. (2011). Reengineering Translational Science: The Time is Right. *Science Translational Medicine*. Retrieved from http://stm.sciencemag.org/content/3/90/90cm17.short

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9397339

Conover, C. (2012). Health Care Regulation: A $169 Billion Hidden Tax. *Cato Instittute*, (527). Retrieved from http://www.cato.org/publications/policy-analysis/health-care-regulation-$169-billion-hidden-tax?print

Cortazar, P., Justice, R., Johnson, J., Sridhara, R., Keegan, P., & Pazdur, R. (2012). US Food and Drug Administration approval overview in metastatic breast cancer. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *30*(14), 1705–11. http://doi.org/10.1200/JCO.2011.39.2613

Crossing the Quality Chasm: A New Health System for the 21st Century. (2001). *Institute of Medicine*.

Cutler, D. M., & Ghosh, K. (2012). The potential for cost savings through bundled episode payments. *The New England Journal of Medicine*, *366*(12), 1075–7. http://doi.org/10.1056/NEJMp1113361

Dabbs, D. J., Klein, M. E., Mohsin, S. K., Tubbs, R. R., Shuai, Y., & Bhargava, R. (2011). High false-negative rate of HER2 quantitative reverse transcription polymerase chain reaction of the Oncotype DX test: an independent quality assurance study. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *29*(32), 4279–85. http://doi.org/10.1200/JCO.2011.34.7963

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 Moor, J. S., Mariotto, A. B., Parry, C., Alfano, C. M., Padgett, L., Kent, E. E., … Rowland, J. H. (2013). Cancer survivors in the United States: prevalence across the survivorship trajectory and implications for care. *Cancer Epidemiology, Biomarkers & Prevention : A Publication of the American Association for Cancer Research, Cosponsored by the American Society of Preventive Oncology*, *22*(4), 561–70. http://doi.org/10.1158/1055-9965.EPI-12-1356

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3386438

DeSantis, C., Ma, J., Bryan, L., & Jemal, A. (2014). Breast cancer statistics, 2013. *CA: A Cancer Journal for Clinicians*, *64*(1), 52–62. http://doi.org/10.3322/caac.21203

Domchek, S. M., Bradbury, A., Garber, J. E., Offit, K., & Robson, M. E. (2013). Multiplex genetic testing for cancer susceptibility: out on the high wire without a net? *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *31*(10), 1267–70. http://doi.org/10.1200/JCO.2012.46.9403

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. Retrieved from http://www.aaai.org/AITopics/classic/Buchanan/Buchanan24.pdf

Farmer, S. A., Black, B., & Bonow, R. O. (2013). Tension between quality measurement, public quality reporting, and pay for performance. *JAMA : The Journal of the American Medical Association*, *309*(4), 349–50. http://doi.org/10.1001/jama.2012.191276

Flanagan, M. B., Dabbs, D. J., Brufsky, A. M., Beriwal, S., & Bhargava, R. (2008). Histopathologic variables predict Oncotype DX recurrence score. *Modern Pathology : An Official Journal of the United States and Canadian Academy of Pathology, Inc*, *21*(10), 1255–61. http://doi.org/10.1038/modpathol.2008.54

Foulkes, W. D., Smith, I. E., & Reis-Filho, J. S. (2010). Triple-negative breast cancer. *The New England Journal of Medicine*, *363*(20), 1938–48. http://doi.org/10.1056/NEJMra1001389

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

Fuchs, V. R. (2013). The gross domestic product and health care spending. *The New England Journal of Medicine*, *369*(2), 107–9. http://doi.org/10.1056/NEJMp1305298

Google. (2014). Google Charts.

Green, E. D., & Guyer, M. S. (2011). Charting a course for genomic medicine from base pairs to bedside. *Nature*, *470*(7333), 204–13. http://doi.org/10.1038/nature09764

Green, L. W., Ottoson, J. M., García, C., & Hiatt, R. A. (2009). Diffusion theory and knowledge dissemination, utilization, and integration in public health. *Annual Review of Public Health*, *30*, 151–74. http://doi.org/10.1146/annurev.publhealth.031308.100049

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

Grossmann, C., Sanders, J., & English, R. A. (2013). *Large Simple Trials and Knowledge Generation in a Learning Health System Workshop Summary*.

Hahsler, M., & Chelluboina, S. (2011). The arules R-Package Ecosystem : Analyzing Interesting Patterns from Large Transaction Data Sets. *The Journal of Machine …*, *12*, 1977–1981. Retrieved from http://dl.acm.org/citation.cfm?id=2021064

Haimowitz, I. J., & Kohane, I. S. (1993). An epistemology for clinically significant trends, 176–181. Retrieved from http://dl.acm.org/citation.cfm?id=1867270.1867297

Haimowitz, I., & Kohane, I. (1993). Automated trend detection with alternate temporal hypotheses. *IJCAI*, 146–151. Retrieved from http://ijcai.org/Past Proceedings/IJCAI-93-VOL1/PDF/021.pdf

Hamburg, M. A., & Collins, F. S. (2010). The path to personalized medicine. *The New England Journal of Medicine*, *363*(4), 301–4. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/20551152

Haut, E. R., & Pronovost, P. J. (2011). Surveillance bias in outcomes reporting. *JAMA : The Journal of the American Medical Association*, *305*(23), 2462–3. http://doi.org/10.1001/jama.2011.822

Hewitt, M., & Simone, J. V. (1999). *Ensuring Quality Cancer Care*. Institute of Medicine.

Higgins, A., Veselovskiy, G., & McKown, L. (2013). Provider performance measures in private and public programs: achieving meaningful alignment with flexibility to innovate. *Health Affairs (Project Hope)*, *32*(8), 1453–61. http://doi.org/10.1377/hlthaff.2013.0007

Higgins, M. J., & Baselga, J. (2011). Targeted therapies for breast cancer. *The Journal of Clinical Investigation*, *121*(10), 3797–803. http://doi.org/10.1172/JCI57152

Howie, L., Hirsch, B., Locklear, T., & Abernethy, A. P. (2014). Assessing the value of patient-generated data to comparative effectiveness research. *Health Affairs (Project Hope)*, *33*(7), 1220–8. http://doi.org/10.1377/hlthaff.2014.0225

IBM Netezza Data Warehouse Appliances – The Simple Data Warehouse Appliance for Serious Analytics. (2014, April 1). IBM Corporation. Retrieved from http://www-01.ibm.com/software/data/netezza/

IOM, I. of M. (2000). *To err is human: building a safer health system*. (L. Kohn, J. Corrigan, & M. Donaldson, Eds.)*National Academy Press*. Washington D.C.: National Academy Press.

Jacobson, M., Earle, C. C., Price, M., & Newhouse, J. P. (2010). How Medicare's payment cuts for cancer chemotherapy drugs changed patterns of treatment. *Health Affairs (Project Hope)*, *29*(7), 1391–9. http://doi.org/10.1377/hlthaff.2009.0563

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

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1766328

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1943788

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2005832

Kahn, M., Tu, S., & Fagan, L. (1991). TQuery: A Context-Sensitive temporal query language. *Computers and Biomedical Research*. Retrieved from http://www.sciencedirect.com/science/article/pii/001048099190016P

King, S. (2001). An All-Consuming Cause: Breast Cancer, Corporate Philanthropy, and the Market for Generosity. *Social Text*, *19*(4), 115–143. Retrieved from http://muse.jhu.edu/journals/social_text/v019/19.4king.html

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

Kohane, I., & Haimowitz, I. (1993). Hypothesis-driven data abstraction with trend templates. … *of the Annual Symposium on Computer …*. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2248548/

Kohn, L. T., Corrigan, J. M., & Molla, S. (2000). *To Err Is Human: Building a Safer Health System*.

Kolender, E. (2009). Picture perfect staffing in today's economy. *Journal of Registry Management*, 153–154. Retrieved from http://www.ncra-usa.org/files/public/Winter2009_V36.4web.pdf#page=53

Korah, M. P., Sener, S. F., & Tripathy, D. (2012). Implications of Omitting Radiation After Breast Conserving Surgery in Elderly Women With Low-risk Invasive Breast Cancer. *International Journal of Radiation Oncology*Biology*Physics*, *84*(3), S34. http://doi.org/10.1016/j.ijrobp.2012.07.094

Kowalski, R., & Sergot, M. (1989). A logic-based calculus of events. *Foundations of Knowledge Base Management*, *4*(June 1975). Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-83397-7_2

Liedtke, C., Mazouni, C., Hess, K. R., André, F., Tordai, A., Mejia, J. a, … Pusztai, L. (2008). Response to neoadjuvant therapy and long-term survival in patients with triple-negative breast cancer. *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, *26*(8), 1275–81. http://doi.org/10.1200/JCO.2007.14.4147

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

McCahill, L. E., Single, R. M., Aiello Bowles, E. J., Feigelson, H. S., James, T. a, Barney, T., … Onitilo, A. a. (2012). Variability in reexcision following breast conservation surgery. *JAMA*, *307*(5), 467–75. http://doi.org/10.1001/jama.2012.43

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

McGlynn, E. A., Asch, S. M., Adams, J., Keesey, J., Hicks, J., DeCristofaro, A., & Kerr, E. A. (2003). The quality of health care delivered to adults in the United States. *The New England Journal of Medicine*, *348*(26), 2635–45. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12826639

Meldrum, C., Doyle, M. A., & Tothill, R. W. (2011). Next-generation sequencing for cancer diagnostics: a practical perspective. *The Clinical Biochemist. Reviews / Australian Association of Clinical Biochemists*, *32*(4), 177–95. Retrieved from http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3219767&tool=pmcentrez&re

ndertype=abstract

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

Meyer, G. S., Nelson, E. C., Pryor, D. B., James, B., Swensen, S. J., Kaplan, G. S., … Hunt, G. C. (2012). More quality measures versus measuring what matters: a call for balance and parsimony. *BMJ Quality & Safety*, *21*(11), 964–8. http://doi.org/10.1136/bmjqs-2012-001081

Meyerson, M., Gabriel, S., & Getz, G. (2010). Advances in understanding cancer genomes through second-generation sequencing. *Nature Reviews. Genetics*, *11*(10), 685–96. http://doi.org/10.1038/nrg2841

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. Retrieved from http://www.ifs.tuwien.ac.at/~silvia/pub/publications/mik_keml97.pdf

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2407929

Morris, Z. S., Wooding, S., & Grant, J. (2011). The answer is 17 years, what is the question: understanding time lags in translational research. *Journal of the Royal Society of Medicine*, *104*(12), 510–20. http://doi.org/10.1258/jrsm.2011.110180

Moses, H., Matheson, D. H. M., Dorsey, E. R., George, B. P., Sadoff, D., & Yoshimura, S. (2013). The anatomy of health care in the United States. *JAMA : The Journal of the American Medical Association*, *310*(18), 1947–63. http://doi.org/10.1001/jama.2013.281425

Moskovitch, R., & Shahar, Y. (2009). Medical Temporal-Knowledge Discovery via Temporal Abstraction. *AMIA Annual Symposium Proceedings*, 452–456. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2815492/

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. Retrieved from http://jamia.bmj.com/content/3/6/367.abstract

Nass, S. J., Patlak, M., & Forum, C. P. (2013). *Implementing a National Cancer Clinical Trials System for the 21st Century : Second Workshop Summary*.

Newcomer, L. N. (2012). Changing physician incentives for cancer care to reward better patient outcomes instead of use of more costly drugs. *Health Affairs (Project Hope)*, *31*(4), 780–5.

http://doi.org/10.1377/hlthaff.2012.0002

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*.

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. Retrieved from http://link.springer.com/article/10.1023/A:1008711006591

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

Olopade, O. I., Grushko, T. A., Nanda, R., & Huo, D. (2008). Advances in breast cancer: pathways to personalized medicine. *Clinical Cancer Research : An Official Journal of the American Association for Cancer Research*, *14*(24), 7988–99. http://doi.org/10.1158/1078-0432.CCR-08-1211

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

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–ff. 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*. Retrieved from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Graphical+Summary+of+Patient+Status#0

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

R Core Team. (2013). R: A Language and Environment for Statistical Computing.

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. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2655843/

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

SB, E., DR, B., & CC, C. (2010). AJCC (American Joint Committee on Cancer). Cancer Staging Manual, 7th edition. Retrieved October 10, 2014, from http://web2.facs.org/cstage0205/breast/Breastschema.html

Schmidt, M. (2008). The Sankey Diagram in Energy and Material Flow Management. *Journal of Industrial Ecology*, *12*(1), 82–94. http://doi.org/10.1111/j.1530-9290.2008.00004.x

Schneider, E. C., Hussey, P. S., & Schnyer, C. (2011). Payment Reform: Analysis of Models and Performance Measurement Implications. *RAND Corporation*. Retrieved from http://www.rand.org/pubs/technical_reports/TR841.html

Shahar, Y. (2013). The Elicitation, Representation, Application, and Automated Discovery of Time-Oriented Declarative Clinical Knowledge. *Process Support and Knowledge Representation in …*, 1–29. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-36438-9_1

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. Retrieved from http://www.sciencedirect.com/science/article/pii/S0933365798000153

Shahar, Y., & Musen, M. (1992). RESUME: A Temporal-Abstraction System for Patient Monitoring. *Proceedings of the Annual Symposium on …*. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2248062/

Shahar, Y., & Musen, M. A. (1996). Knowledge-based temporal abstraction in clinical domains. *Artificial Intelligence in Medicine*, *8*(3), 267–298. Retrieved from http://dx.doi.org/10.1016/0933-3657(95)00036-4

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

http://www.sciencedirect.com/science/article/pii/104281439290027X

Shahian, D. M., Wolf, R. E., Iezzoni, L. I., Kirle, L., & Normand, S.-L. T. (2010). Variability in the measurement of hospital-wide mortality rates. *The New England Journal of Medicine*, *363*(26), 2530–9. http://doi.org/10.1056/NEJMsa1006396

Sharf, B. F. (2001). Out Of The Closet And Into The Legislature: Breast Cancer Stories. *Health Affairs*, *20*(1), 213–218. http://doi.org/10.1377/hlthaff.20.1.213

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. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2464549/

Shrank, W. (2013). The Center For Medicare And Medicaid Innovation's blueprint for rapid-cycle evaluation of new care and payment models. *Health Affairs (Project Hope)*, *32*(4), 807–12. http://doi.org/10.1377/hlthaff.2013.0216

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.

Smart, C. R., Hartmann, W. H., Beahrs, O. H., & Garfinkel, L. (1993). Insights into breast cancer screening of younger women. Evidence from the 14-year follow-up of the Breast Cancer Detection Demonstration Project. *Cancer*, *72*(4 Suppl), 1449–56. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8339237

Smith, T. J., & Hillner, B. E. (2011). Bending the cost curve in cancer care. *The New England Journal of Medicine*, *364*(21), 2060–5. http://doi.org/10.1056/NEJMsb1013826

Spence, I. (2006). William Playfair and the psychology of graphs. *2006 JSM Proceedings, American Statistical …*, 2426–2436. Retrieved from https://noppa.aalto.fi/noppa/kurssi/25333/materiaali/25333_media_lab_2.pdf

Spinks, T. E., Walters, R., Feeley, T. W., Albright, H. W., Jordan, V. S., Bingham, J., & Burke, T. W. (2011). Improving cancer care through public reporting of meaningful quality measures. *Health Affairs (Project Hope)*, *30*(4), 664–72. http://doi.org/10.1377/hlthaff.2011.0089

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 …*. Retrieved from http://www.imia-medinfo.org/wg_idamap/idamap/idamap2003/Spokoiny.pdf

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

Stockdale, H., & Guillory, K. (2013). *Why cancer patients depend on medicare for critical coverage*. Retrieved from http://www.acscan.org/content/wp-content/uploads/2013/06/2013-Medicare-Chartbook-Online-Version.pdf

Sullivan, R., Peppercorn, J., Sikora, K., Zalcberg, J., Meropol, N. J., Amir, E., … Aapro, M. (2011). Delivering affordable cancer care in high-income countries. *The Lancet. Oncology*, *12*(10), 933–80. http://doi.org/10.1016/S1470-2045(11)70141-3

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9762927

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

Tangka, F., Subramanian, S., Beebe, M. C., Trebino, D., & Michaud, F. (2010). Economic assessment of central cancer registry operations, Part III: Results from 5 programs. *Journal of Registry Management*, *37*(4), 152–5. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/21688745

Targeted Therapies for Breast Cancer Tutorial - National Cancer Institute. (n.d.). Retrieved September 27, 2014, from http://www.cancer.gov/cancertopics/understandingcancer/targetedtherapies/breastcancer_htmlcourse/page7

The Perl Programming Language - www.perl.org. (n.d.). Retrieved June 28, 2015, from https://www.perl.org/

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

Vanneman, M., & Dranoff, G. (2012). Combining immunotherapy and targeted therapies in cancer treatment. *Nature Reviews. Cancer*, *12*(4), 237–51. http://doi.org/10.1038/nrc3237

Vinh-Hung, V., & Verschraegen, C. (2004). Breast-Conserving Surgery With or Without Radiotherapy: Pooled-Analysis for Risks of Ipsilateral Breast Tumor Recurrence and Mortality. *JNCI Journal of the National Cancer Institute*, *96*(2), 115–121. http://doi.org/10.1093/jnci/djh013

Vogelstein, B., Papadopoulos, N., Velculescu, V. E., Zhou, S., Diaz, L. A., & Kinzler, K. W. (2013). Cancer genome landscapes. *Science (New York, N.Y.)*, *339*(6127), 1546–58. http://doi.org/10.1126/science.1235122

Von Minckwitz, G., & Martin, M. (2012). Neoadjuvant treatments for triple-negative breast cancer (TNBC). *Annals of Oncology : Official Journal of the European Society for Medical Oncology / ESMO*, *23 Suppl 6*(Supplement 6), vi35–9. http://doi.org/10.1093/annonc/mds193

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

Winchester, D. P., Stewart, A. K., Phillips, J. L., & Ward, E. E. (2010). The national cancer data base: past, present, and future. *Annals of Surgical Oncology*, *17*(1), 4–7. http://doi.org/10.1245/s10434-009-0771-3

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. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12749723

Wongsuphasawat, K. (2011). LifeFlow: visualizing an overview of event sequences. *Proceedings of the …*, 1747–1756. Retrieved from http://dl.acm.org/citation.cfm?id=1979196

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

Zaki, M. (2001). SPADE: An efficient algorithm for mining frequent sequences. *Machine Learning*, 31–60. Retrieved from http://link.springer.com/article/10.1023/A:1007652502315

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. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6523038