Toward Data-Driven Learning Healthcare Systems in Interventional Radiology: Implementation to Evaluate Venous Stent Patency

David M. Cohn¹ · Tarub S. Mabud¹ · Victoria A. Arendt² · Andre D. Souffrant¹ · Gyeong S. Jeon³ · Xiao An⁴ · William T. Kuo¹ · Daniel Y. Sze¹ · Lawrence V. Hofmann¹ · Daniel L. Rubin¹

Published online: 24 October 2019

© Society for Imaging Informatics in Medicine 2019

Abstract

We developed a code and data-driven system (learning healthcare system) for gleaning actionable clinical insight from interventional radiology (IR) data. To this end, we constructed a workflow for the collection, processing and analysis of electronic health record (EHR), imaging, and cancer registry data for a cohort of interventional radiology patients seen in the IR Clinic at our institution over a more than 20-year period. As part of this pipeline, we created a database in REDCap (VITAL) to store raw data, as collected by a team of clinical investigators and the Data Coordinating Center at our university. We developed a single, universal pre-processing codebank for our VITAL data in R; in addition, we also wrote widely extendable and easily modifiable analysis code in R that presents results from summary statistics, statistical tests, visualizations, Kaplan-Meier analyses, and Cox proportional hazard modeling, among other analysis techniques. We present our findings for a test case of supra versus infra-inguinal ligament stenting. The developed pre-processing and analysis pipelines were memory and speed-efficient, with both pipelines running in less than 2 min. Three different supra-inguinal ligament veins had a statistically significant improvement in vein diameters post-stenting versus pre-stenting, while no infra-inguinal ligament veins had a statistically significant improvement (due either to an insufficient sample size or a non-significant p value). However, infra-inguinal ligament stenting was not associated with worse restenosis or patency outcomes in either a univariate (summary-statistics and Kaplan-Meier based) or multivariate (Cox proportional hazard model based) analysis.

Keywords Data mining · Learning healthcare system · Interventional radiology · Inguinal ligament

Introduction

The advent of EHR is producing an unprecedented amount of medical data that can be leveraged in analyses to improve clinical decision-making. Systems built using this paradigm are called “learning healthcare systems.” However, a significant portion of EHR data is poorly structured, and there are technical complexities in collecting, managing, and analyzing this data, which hinders the development of such systems. In IR, these difficulties are further complicated by the need to

assess radiographic data, which typically reside in separate systems from institutional EHRs and require special software and skill-sets to annotate.

Clinical informatics and data mining techniques that are increasingly prevalent in a variety of clinical domains are underutilized in IR. Traditional research studies in IR literature tend to involve small cohorts of hand-curated patient populations, for which data is costly to collect, and whose results may not generalize to larger populations. In this study, we describe our framework for the capture of comprehensive retrospective patient-level medical data into a structured, longitudinal database for over 20 years of IR patients: the Venous InTerventionAL (VITAL) database. To demonstrate the capability of this resource to enable the learning healthcare paradigm, we apply this database and the analytical tools we built to investigate the safety and efficacy of a particular IR therapy for chronic deep venous thrombosis: venous stenting across the inguinal ligament.

Arterial stenting across joints is associated with well-characterized risks of stent damage and neointimal hyperplasia,

IRB: Approved by Stanford IRB, VITAL #33192

✉ Daniel L. Rubin
dlrubin@stanford.edu

¹ Stanford University School of Medicine, Stanford, CA, USA

² Cleveland Clinic, Cleveland, OH, USA

³ CHA University Bundang Medical Center, Seongnam, South Korea

⁴ Shanghai General Hospital, Shanghai, China

leading to lower patency rates. [1–3] Venous stents placed across joints for the treatment of chronic outflow obstruction are hypothesized to experience an analogous risk to long-term patency. In particular, there is considerable reluctance among IR practitioners to place stents that cross the inguinal ligament, even when the veins crossing the inguinal ligament are heavily involved in thrombotic or other occlusive disease processes, due to fear of stent compression, re-thrombosis, or damage. [4] Prior investigations of venous stenting across the inguinal ligament have corroborated concerns regarding poor outcomes. A study of 493 limbs with venous stents terminating either supra versus infra-inguinal ligament found significantly worse secondary patency rates among the infra-inguinal cohort compared to the supra-inguinal cohort (84% vs 90%; $p < 0.001$) for patients with post-thrombotic occlusion. [4] Similarly, an abstract describing 94 patients with post-thrombotic venous occlusion found that stents placed across the inguinal ligament had significantly worse secondary patency rates after 1 year than those placed supra-inguinal ligament (82% vs 100%; $p < 0.05$). [5]

Despite these findings, the authors of both studies contend that stenting across the inguinal ligament is safe and effective, given the relatively high, albeit significantly lower, patency associated with stents that cross the inguinal ligament. However, neither of these studies considered a comprehensive set of important covariates and potential confounders in their analyses of longitudinal patency rates. Thus, while limited data exists for infra-inguinal stenting, they do not fully support assertions of long-term efficacy, and they have not engendered enough confidence among many IR radiologists to pursue infra-inguinal ligament stenting.

The VITAL database was built from a large collection of curated medical data, which contains detailed individual-level clinical, procedural, imaging, and longitudinal follow-up data for patients who received IR interventions at our institution over the past two decades. This resource was designed for creating learning healthcare systems in IR and contains an extensive set of covariates, which include potential confounders and follow-up indicators that characterize complications and long-term outcomes. We show the value of building such a resource for learning healthcare systems in IR by leveraging VITAL to investigate supra versus infra-inguinal ligament stenting.

Materials and Methods

VITAL Database Construction

In constructing the VITAL database, we began with a cohort of 537 patients that received IR venous procedures between July 1996 and April 2018, procured using a cohort selection tool at our institution (STARR) that identifies patients based on diagnosis and procedural codes. [6] Subsequently, the data set was ultimately extended to include 684 patients with relevant IR data.

The VITAL database was built using the REDCap software platform [7] by extracting 816 clinical predictors derived from each patient's medical and imaging record, including patient demographics, medical history, venous thromboembolism (VTE) and risk factors, IR clinic visits, procedures, non-IR encounters, imaging, medications, labs, and cancer registry diagnosis data (Fig. 1). REDCap is a secure, web-based application that supports data capture for research studies, providing (1) an intuitive interface for validated data entry, (2) audit trails for tracking data manipulation and export procedures, (3) automated export procedures for data downloads to statistical packages, and (4) procedures for importing data from external sources. [7]

Figure 1 shows the flow of data from patient records to the VITAL database. The variables derived from a patient's EHR and imaging records (EPIC and PACS, respectively) were collected by a team of investigators by acquiring data from the clinical data system (STARR) and our institution's imaging system (PACS), respectively. We procured structured data on patient demographics, medical history, VTE and risk factors, IR clinic visits, and non-IR encounters from the EHR record. Imaging-related data, such as measurement data of deep vein thrombosis (DVT) present, stents and vessels visualized, was obtained from the PACS imaging records. Finally, the Data Coordinating Center at our institution pulled cancer registry, medications, labs, last follow-up, and mortality outcome data by using procedure and billing codes, keywords, records, and clinical notes taken from EPIC, our institution's cancer registry and the Social Security Death Index (SSDI), among other sources. [8]

Ninety-seven of the 816 predictors in the REDCap database were acquired by the DCC; in turn, 119 variables make up the imaging predictors derived via PACS. Finally, the remaining predictors were obtained via EHR. In total, of the 558,144 data elements (684 patients $\times$ 816 predictors) in VITAL, 66,348 data elements (684 patients $\times$ 97 predictors) were sourced from the DCC, 81,396 imaging data elements (684 patients $\times$ 119 predictors) were procured from PACS, and the remaining 410,400 data elements were EHR-based variables. Examples of data derived by the DCC include cancer staging and treatment (chemotherapy and radiation) data; in turn, vein diameter and vein-level restenosis data were among the variables collected from imaging studies. Finally, the number/location of stents and angioplasty balloons used for IR procedures were among the several hundred predictors sourced from the EHR.

Once structured data had been obtained using the three mechanisms described above, the structured variables were entered into the VITAL database by a team of investigators using the user interface available in REDCap. This interface, as represented by part of the procedure data entry form shown in Fig. 2, enables the input of both structured and free-text in a variety of possible formats (dropdown, radio buttons, free

Fig. 1 Structured data is derived using a combination of clinician evaluation (EHR and imaging) and a data acquisition workflow (Data Coordinating Center Data Acquisition Pipeline) and then inputted into the VITAL database using the REDCap user interface (GUI); figure created in Adobe Photoshop CC

text, date, etc.); free-text can be entered using provided text boxes (Full Procedure Report). Categorical variables can be captured in either radio button (Type of Note) or check box (Access Site) format, depending on whether the categorical variable can take on a single value or multiple values concurrently. Finally, date-times can be inputted in special text boxes (Procedure Date) that constrain the date-time elements to a pre-determined format (such as MM-DD-YY etc.). The flexibility in possible data formats available in REDCap, coupled with the organized and systematic structures associated with these formats, and an easy-to-follow visual interface, enabled the population of the VITAL database.

Pre-processing Workflow

We developed a data pre-processing and cleaning pipeline to process the raw structured data from REDCap. This pipeline converts numerical encodings of categorical variables (as

stored in REDCap) to their corresponding string-based names, and parses measurement entries (such as for stent lengths and diameters) to distinguish measurement values from their corresponding units. The pipeline also standardizes string-based names to a consistent format (such as capitalizing the first letter of every word in a stent brand name).

For inputs with multiple possible data values (such as multiple listed dosages for anticoagulation plan data), our code cleans the data by identifying the correct value amid the multiple values, based on identified data collection trends. For time-duration data, such as the length of anticoagulation treatment, the code identifies the unit of time used, as well as the numerical duration value, and standardizes the duration measurements based on a consistent unit of measurement. The code also removes improper date-time entries and standardizes both complete and incomplete date-time entries by first identifying the original date-time format, filling in any potential missing date-time components (such as a missing month

Procedure

Record ID

Introduction

Procedure Date

Time Since Previous Procedure

Should this count as a reintervention?

- ☐ Yes
☐ No
 (i.e. was it unplanned?)

Attending

Fellow

Full Procedure Report

Type of note

- ☐ Parseable
☐ Free Text

Procedure Performed

On which veins was this procedure performed?

- ☐ SVC and above
☐ IVC and below
☐ Other

Access Site

- ☐ R Internal Jugular Vein
☐ R Common Femoral Vein
☐ R (Superficial) Femoral Vein
☐ R Popliteal Vein
☐ R Ankle Vein
☐ R Greater Saphenous Vein
☐ R Small Saphenous Vein
☐ L Internal Jugular Vein
☐ L Common Femoral Vein
☐ L (Superficial) Femoral Vein
☐ L Popliteal Vein
☐ L Ankle Vein
☐ L Greater Saphenous Vein
☐ L Small Saphenous Vein

Fig. 2 Data entry form in REDCap

or day) based on a consistent set of assumptions regarding the context of date (i.e., does the date represent an IR clinic visit date versus the date of a procedure in the medical history versus the date of the emergence of a risk factor etc.), before standardizing the appearance of the date-time value.

For certain numerical values (such as stent lengths and diameters) with a pre-defined range of possible values that

one would expect in clinic, we identified outliers (such as a stent diameter less than 1 cm or a stent diameter greater than 40 cm) through thresholding, with manual correction of the values in REDCap after additional confirmation by a member of the data collection team. We also cleaned potentially noisy DVT and restenosis imaging data by examining each patient's imaging follow-up in R and identifying incongruencies (such

as the appearance of in-stent restenosis prior the placement of any stent etc.)

Finally, our pre-processing codebank creates more compact and memory-efficient units of data storage by dividing the sections of the VITAL database described above (Demographics, Medical History, VTE etc.) into further sub-sections (for example, dividing medical history into thrombophilia, venous medical history, and other medical history) and removing rows with missing or not available data elements from these storage dataframes. These compact representations allow for more efficient analysis of the data in the subsequent analysis pipeline.

The pre-processing pipeline was primarily constructed using the dplyr [9] and tidyr [10] packages in R [11] (R-equivalent of pandas package in Python). dplyr and tidyr are memory and speed-efficient packages that are standards for data wrangling in R; they enable the automatic calculation of new variables based on the input of previously collected or calculated predictors, the filtering of data based on a set of value(s) or threshold, and even the merging of data sets on a common variable without having to manually perform these operations. Furthermore, dplyr and tidyr allow for the sequential performance of multiple operations automatically through use of an operator, without having to make function calls one at a time. In this way, data can be instantaneously processed and cleaned, while dramatically reducing the risk of human error. We made our pre-processing codebank publicly available on GitHub. [12]

Data Analysis Workflow

Our data analysis pipeline sequentially integrates with our pre-processing workflow, as the first method called in the analysis pipeline is our pre-processing package. Our data analysis code is also publicly available on GitHub. [13] The data analysis codebase calculates summary statistics of variables, performs statistical tests (such as comparing two cohorts of patients via two-sample statistical test methodologies such as *t* tests, Chi-square tests etc.), conducts Cox proportional hazard modeling and Kaplan-Meier analyses based on selected covariates, and derives data visualizations (Table 1). Summary statistics are

Table 1 Major data analysis modules and primary R package(s) used in module construction

Module name	Primary package(s) used
Patency	dplyr
Summary statistics	dplyr
Statistical tests	stats
Visualizations	ggplot2
Kaplan-Meier	Survival, Survminer
Cox proportional hazard	Survival

primarily generated using the summary method in the dplyr package. Statistical tests are conducted using the stats package in R.

We perform Kaplan-Meier analyses and Cox proportional hazard modeling using the Survival package in R. [14] While Kaplan-Meier analyses and Cox proportional hazard modeling were originally developed for performing survival analyses, namely time-to-mortality analyses, they are nevertheless generalizable to other time-to-event analyses [15], such as the time-until-restenosis and time-until-primary-patency-lost analyses that we used in our demonstrated use case for this study. The benefit of utilizing Kaplan-Meier and Cox proportional hazard modeling, in addition to basic event-based summary statistics, is that these methods take into account patient censoring/lost to follow-up in calculating risk/event probabilities, whereas simple summary statistics methods will not.

Our code module reports 95% confidence intervals for our Kaplan-Meier plots and assesses differences in Kaplan-Meier curves using a log-rank test. [16] We also display hazard ratios with corresponding *p* values for the Cox proportional hazard modeling using the summary function in R.

Our analysis code creates visualizations using the ggplot2 package in R. [17] ggplot2 enables the creation of a vast array of possible graphs and diagrams through a universal command, with only a small additional add-on flag to specify the visualization type. In addition, ggplot2, in concert with other packages in R, also allows for the automatic creation of figures, based on one or more graphs, without manual construction. Thus, any modifications of the code pipeline upstream will automatically change the figure, without having to create a new figure manually from scratch.

Our pre-processing and analysis code was designed to be run in RStudio Desktop [18], which is a widely used integrated development environment (IDE) for running R code. RStudio Desktop allows code to be run either line-by-line, or script-by-script, by highlighting the code to run and pressing Run. In addition, dataframes can be visually examined by clicking on the appropriate data frame in the Environment section of the IDE. Finally, visualization plots can be displayed via the Plots tab in the bottom-right corner of the display. In this way, code, data structures and plots can all be examined and displayed in one integrated screen.

Our pre-processing and analysis workflows, along with the dataframe output and a Kaplan-Meier visualization, are shown in RStudio in Fig. 3. In turn, a diagram showing the sequence of code making up our overall workflow is shown in Fig. 4.

Demonstration Use Case: Supra Versus Infra-Inguinal Ligament Cohort Construction and Variable Calculation

To demonstrate the utility of our system, we used our code to explore a specific clinical question: whether infra-inguinal

Fig. 3 Learning healthcare (pre-processing and analysis) system in RStudio Desktop, with its outputted results in the form of dataframes and visualizations (Kaplan-Meier figure shown); figure created in R/RStudio

ligament stenting led to significantly worse patient outcomes in vein diameters, restenosis, and patency. To this end, we used our pipeline to generate summary statistics, visualizations, and statistical tests to assess changes in vein diameters pre-stenting versus post-stenting, while we used summary statistics, statistical tests, Kaplan-Meier analyses, and Cox proportional hazard modeling to assess restenosis and patency outcomes (Table 1). We describe each of these analyses in additional detail below.

To acquire the cohort of patients used in our analysis, we applied the filter method in the dplyr package in R to remove patients that had documented venous stenting at an outside hospital prior to their first venous stenting procedure at our institution, or who did not have either left or right leg lower extremity stenting during their first stenting procedure at our university, or who had discontinuous inferior vena cava (IVC) stenting and either left or right leg femoral stenting. This filter operation resulted in a total of 344 patients and 448 limbs; 340 limbs had supra-inguinal ligament stenting, while 108 limbs had infra-inguinal ligament stenting.

Vein Diameters

Our vein diameter code preferentially selected vein diameter measurements that were procured using either computed tomography (CT) or magnetic resonance (MR) venography imaging within 1 year of the stenting procedure, with post stenting measurements taking place prior to any loss of primary patency status/re-intervention (if applicable).

Restenosis and Patency

We utilized our Cox proportional hazard module to determine statistically significant factors associated with restenosis and patency outcomes. We calculated or derived the following variables from our VITAL database: supra versus infra-inguinal ligament stenting status, the gender, ethnicity, and age of the patient at the time of the stenting intervention, whether IVC stenting was also performed at the time of the stenting intervention (Yes or No), the thrombophilia status of the patient, the laterality of the limb, and the patients most recent DVT/Lymphedema status in the limb prior to the stenting intervention (DVT, lymphedema, both or none; whether the DVT was provoked or not was also incorporated into this categorical factor).

Stents were categorized as supra- or infra-inguinal using a dataframe of stented veins produced by our pre-processing codebank. Limbs were classified as having infra-inguinal stenting if the limb had at least one stent placed in the common femoral, femoral, profunda femoris, or popliteal veins, irrespective of proximal stent placement, during the index procedure at the study institution. Limbs were classified as having supra-inguinal stenting if the limb had at least one stent placed in the common iliac or external iliac veins (irrespective of inferior vena cava stenting), and no stents distal to the external iliac vein, during the index procedure at the study institution. Patients with exclusive inferior vena cava stent placement were excluded from the cohort, because patients with IVC stenting only typically represent a distinct IR patient cohort with a different

Fig. 4 Pre-processing and analysis workflow diagram with significant modules indicated; figure created in Adobe Photoshop CC

diagnosis and severity of disease. The laterality of the limb was also calculated from the dataframe of stented veins, based on the laterality of the vein indicated in the vein name.

To calculate IVC status, we extracted a table for our supra versus infra-inguinal ligament cohort containing a unique patient identifier, as well as a unique procedure identifier, and filtered our stented vein dataframe on these combinations of patient and procedure identifiers; subsequently, we preferentially selected and identified veins from the remaining dataframe that represented suprarenal IVC and/or infrarenal IVC veins. The most recent DVT and/or lymphedema for each limb was calculated by first filtering our dataframe of DVT and lymphedema events produced by the pre-processing pipeline on the laterality of the limb and the record id of the patient. Next, we calculated the span between the remaining DVT and/or lymphedema events and the patients' first intervention at

our institution in R and selected the event with the smallest span between dates.

With regards to patency for our supra versus infra-inguinal ligament experiments, primary patency is defined to have begun on the follow-up start date, namely the date of each patients' first venous stenting procedure at our institution. Primary patency is assumed to be maintained until a major thrombotic complication (in-stent rethrombosis, DVT, or lymphedema) was reported, or the patient required one of the following venous interventions, in either the IVC or the limb in question, to maintain patency: thrombolysis, thrombectomy, recanalization, stenting, angioplasty. We base this primary patency calculation on accepted interventional radiology reporting standards. [19] Patency calculations were performed by our patency module, which utilizes dplyr, among other R packages, to implement the clinical patency criterion described above.

Results

VITAL Database and Learning Healthcare Development Time and Runtime Results

VITAL database design took ~ 50 man-hours to complete, while the non-imaging variables in VITAL took ~ 510 man-hours to collect or ~ 45 min per case. The collection of imaging variables in VITAL took between ~ 2000 and ~ 4000 man-hours to finish. The initial draft of the pre-processing workflow took 80 man-hours to code and takes ~15 s to run in its entirety. In turn, the data analysis pipeline was written in another 80 man-hours and takes ~ 1 min in runtime to complete all requisite calculations.

Results from Pursuing Example Use Case

Vein Diameters

Our analysis pipeline generated visualizations (Visualizations Module) highlighting the change in maximum vein diameters pre-stenting versus post-stenting. For all veins, the average maximum diameter measurement was larger after stenting, as opposed to before stenting. However, to examine whether these improvements were statistically significant, we subsequently considered cases with both pre-stent and post-stent measurements as part of a paired *t* test (statistical test module).

For veins with sufficient sample sizes, our queries and analyses showed that the difference between pre-stenting and post-stenting diameters for RCIV (sample size: 34, $p = 0.000418$), LEIV (sample size: 41, $p = 0.000543$), and LCIV (sample size: 70, $p = 1.83 \times 10^{-8}$) was strongly statistically significant, while the improvement pre-stenting versus post-stenting was nearly statistically significant for LCFV (sample size: 20, $p = 0.0802$). There was no statistically significant difference between pre-stenting and post-stenting measurements for REIV (sample size: 20, $p = 0.395$).

Restenosis

Using our Summary Statistics component, we determined that 66 out of the 340 limbs with supra-inguinal ligament stenting developed restenosis in at least one lower-extremity stented vein during follow-up (19.4%). In turn, 27 out of 108 limbs with infra-inguinal ligament stenting developed restenosis in at least one stented vein during follow-up (25.0%). However, according to our statistical test module, the difference in restenosis rates between the two cohorts was not statistically significantly different (Chi-square test, $p = 0.27$).

Our Kaplan-Meier analysis module computed 5-year post-stenting restenosis probabilities. The risk of a limb developing in-stent restenosis by supra versus infra-inguinal ligament status was not statistically significantly different (log-rank test,

$p = 0.1$; Fig. 5). Nevertheless, the generated Kaplan-Meier curves appear to noticeably diverge after the stenting intervention, with a 5-year restenosis probability for a limb with supra-inguinal ligament stenting of 39.6%, while a limb with infra-inguinal ligament stenting had a 5-year restenosis probability of 50.0%.

Our multivariate limb restenosis Cox proportional hazard model, generated using the Cox proportional hazard module, identified gender to be statistically significantly associated with a higher risk of in-stent restenosis ($p = 0.0097$); supra versus infra-inguinal ligament status was not a statistically significant covariate in the model ($p = 0.24$). Limbs with infra-inguinal ligament stenting had a slightly higher hazard ratio, relative to limbs with supra-inguinal ligament stenting (HR ratio 1.33:1).

Patency

From our Summary Statistics section, we found that 263 out of 340 limbs with supra-inguinal ligament stenting maintained primary patency during the course of follow-up (77.4%). In turn, 90 out of 108 limbs with infra-inguinal ligament stenting maintained primary patency status during follow-up (83.3%). According to our Statistical Test Module, differences in patency rates between the two cohorts were not statistically significantly different (Chi-square test: $p = 0.234$).

Our Kaplan-Meier analysis module showed that, based on 5-year post-stenting primary patency Kaplan-Meier curves, the risk of losing primary patency was not statistically significantly different by supra versus infra-inguinal ligament status (log-rank test: $p = 0.38$; Fig. 6). However, the risk of a limb losing primary patency status with supra-inguinal ligament stenting continues to increase 2+ years after the stenting intervention, with a 5-year primary patency probability of 55.2%; in contrast, the 5-year primary patency probability of a limb with infra-inguinal ligament stenting was 68.7%.

Using our Cox proportional hazard module, we found that IVC stenting at the time of the intervention was associated with a statistically significant increase in the risk of losing primary patency status ($p = 0.0015$). The hazard ratio of IVC stenting to no IVC stenting at the time of the intervention was 2.2:1. There was no statistically significant difference in coefficients/hazard ratios by supra versus infra-inguinal ligament status ($p = 0.432$). Limbs with infra-inguinal ligament stenting were associated with a lower hazard ratio, relative to limbs with supra-inguinal ligament stenting (HR ratio 0.81:1).

Discussion

Using supra versus infra-inguinal ligament stenting as a proof-of-concept experiment, we demonstrate the utility of

Fig. 5 Limb-level restenosis Kaplan-Meier curves for the supra versus infra-inguinal ligament cohorts; figure created in R/RStudio

developing informatics-based learning healthcare systems for data collection, pre-processing, analysis and visualization in order to answer IR research questions. To wit, our use of STARR enabled the efficient identification of our patient cohort, without having to directly access the patient EHR nor write potentially lengthy SQL code. In turn, the use of REDCap for the construction of our database allowed a team of investigators to collaboratively collect a longitudinal database, spanning several hundred patients and predictors, and 20 years of IR data at our institution, without requiring technical expertise in the form of SQL database management. While both STARR and REDCap were previously developed by software developers and informaticians, they nevertheless enabled the creation of a workflow that effectively integrated with our informatics-based pre-processing and data analysis pipeline.

The pre-processing and analysis codebank that we developed in R was able to sift through an extensive database, with hundreds of thousands of data elements, and provide an answer regarding an important question in IR in less than 2 min. The utility of this component of our IR learning healthcare

system also extends to development time; while the planning and construction of the VITAL database took ~ 6 months, the initial construction of the pre-processing code only took ~ 1 week, with subsequent periodic adjustments to the codebank as variables were added to the database, or its database structure modified. Similarly, the drafting of the analysis code only took 2–3 weeks, allowing us to reach a conclusion regarding the question of supra versus infra-inguinal ligament stenting in a relatively short amount of time.

Finally, our learning healthcare system is generalizable to almost any IR venous intervention study at our institution and can be easily modified to account for the collection of new predictors in the future; our pre-processing pipeline can be directly applied, without any additional modifications, to any experiment involving the VITAL database, while our analysis code can be easily modified to meet the needs of subsequent analyses. Furthermore, with platforms such as Github, we can share our codebase with the scientific community and the public writ large, thereby allowing for greater transparency regarding our methodologies, as well as enabling others to provide feedback on our analytical approach more easily.

Fig. 6 Limb-level primary patency Kaplan-Meier Curves for the supra versus infra-inguinal ligament cohorts; figure created in R/RStudio

Our use of a multi-layered analysis approach allowed for different insight regarding the efficacy of infra-inguinal ligament stenting. At the vein level, each iliac and femoral vein cohort, irrespective of laterality, showed improvement in vein diameter post-stenting versus pre-stenting; however, a statistically significant improvement was limited to three of the four iliac veins (RCIV, LEIV, LCIV), while none of the four femoral veins had a statistically significant increase in vein diameter (sample size or otherwise). Furthermore, the increase in vein diameter post-stenting was almost uniformly greater for iliac veins than femoral veins. Taken together, these two findings are suggestive that while femoral vein stenting can be effective, treating iliac veins with stenting may be more efficacious.

This conclusion is further supported by the limb-level restenosis results of our Kaplan-Meier analysis. While the finding was technically not statistically significant ($p = 0.1$), the analysis suggested that limbs with infra-inguinal ligament stenting more frequently developed in-stent restenosis in one or more stented veins. Nevertheless, the limb-level analysis also revealed that limbs with femoral vein stenting do not necessarily require more frequent re-interventions, or are more

susceptible to major thrombotic complications. In fact, infra-inguinal ligament stented limbs more consistently maintained primary patency, as opposed to limbs with supra-inguinal ligament stenting; this finding, according to our Cox proportional hazard model, was likely due to the latter cohort's more frequent association with IVC stenting.

Our study has some limitations. For one, studies such as ours involving the VITAL database are necessarily retrospective, and studies on an independent dataset or in prospective studies would be helpful in confirming our findings. However, retrospective analyses nevertheless enable the identification of interesting hypotheses that can be explored in subsequent studies. An additional limitation is the inherent variability introduced by using data acquired in routine clinical practice, as opposed to data from clinical trials. In clinical practice, multiple biases could be introduced that affect the characteristics of the population (e.g., practice choices inherent in each physician, variations in time for patient follow-up etc.). Thus, validation on independent data is needed. Finally, we applied a very strict definition for in-stent restenosis ($< 5\%$). We have collected more detailed in-stent restenosis percentage data,

binned with the following groupings: 1–25%, 26–50%, 51–75%, 76–99%, and 100% occlusion. This data will enable a more nuanced examination of restenosis outcomes beyond our current binary labels.

In addition, to improve on the data collection time for future work involving IR data, including data that would be pertinent to our VITAL database, a colleague of ours is developing a parser of IR procedure reports, based on a recently implemented standardized EPIC template in the IR department at our institution. This project represents a step forward toward the automated collection of clinical data that nevertheless was not possible for many predictors in this study; however, three noteworthy unresolved challenges toward fully automating the collection of IR data are the lack of standardized formats for many clinical note types, the large inter-clinician variability in radiology measurements such as vein diameters, and the absence of certain important imaging features from standard radiology reports. Nevertheless, the automated parsing of procedure reports represents another step toward a fully automated “learning healthcare” system for data collection, management, analysis, and learning.

Conclusion

In conclusion, we built a prototype learning healthcare system for the collection, pre-processing, and analysis of venous IR data. This system, while applied to the question of supra versus infra-inguinal ligament stenting in this study, is generalizable to a broad array of clinical questions in IR research. In addition, the flexibility of the database architecture, as enabled by REDCap, coupled with a codebase that can be easily modified to reflect changes in the composition of the database, will allow for the possible collection of new clinical variables moving forward. Furthermore, we are currently deploying significant portions of our aforementioned workflow in studies pertaining to IVC atresia, venous stenting and more, thus validating the adaptability of our system. Our work highlights the value of informatics methodology in enabling more efficient and robust IR research. In the end, we hope that our study will serve as a proof-of-concept experiment for the use of informatics techniques in IR studies.

Funding Information This research was funded by the Stanford Medical Scholars Program, as well as the Stanford University Department of Radiology, CHA University Bundang Medical Center and Shanghai General Hospital in the form of salary support for D.M.C., G.S.J., and X.A., respectively.

Compliance with Ethical Standards

L.V.H. is a consultant for Cook Medical and has received royalties and/or equity payouts from Cook Medical, Boston Scientific, Medtronic and Confluent Medical. L.V.H. is also a co-founder, current member of the

Board of Directors, and has equity at Grand Rounds. W.K. is a consultant for Walk Vascular. D.Y.S. is/has been a consultant for the following companies in the last 3 years: Astra-Zeneca, Bayer, BlackSwan Vascular, Boston Scientific, Bristol Myers Squibb, BTG, Eisai, EmbolX, W. L. Gore, Janssen, Koli Medical, RadiAction Medical, Terumo. In addition, D.Y.S. has/has received equity from Confluent Medical and Proteus Digital Health. Finally, D.Y.S. has/has had institutional support in the last 3 years from BTG, W. L. Gore, Merit Medical, and Sirtex.

We received Institutional Review Board approval at our institution. For this retrospective study, formal consent was not required. This article does not contain any studies with animals performed by any of the authors.

This work has been previously published as a pre-print on Research Gate at the following URL:

https://www.researchgate.net/publication/329052256_Toward_Data-Driven_Learning_Healthcare_Systems_in_Interventional_Radiology_Decision_Making_in_Inguinal_Ligament_Stenting

We have made subsequent edits to the paper that was pre-printed to Research Gate to arrive at the paper draft in its current form.

References

1. Babal E, Gilbaran M, Grmen T, Oztürk S: Fracture of popliteal artery stents. *Circ J* 67(7):643645, 2003
2. Andrews RT, Venbrux AC, Magee CA, Bova DA: Placement of a flexible endovascular stent across the femoral joint: An in vivo study in the swine model. *Vasc Interv Radiol* 10(9):12191228, 1999
3. Ballard JL, Sparks SR, Taylor FC, Bergan JJ, Smith DC, Bunt TJ, Killeen JD: Complications of iliac artery stent deployment. *Vasc Surg* 24(4):545555, 1996
4. Neglén P, Tackett TP, Raju S: Venous stenting across the inguinal ligament. *J Vasc Surg* 48(5):12551261, 2008
5. Saha P, Gwozdz A, Hagley D, El-Sayed T, Hunt B, McDonald V, Cohen A, Breen K, Karunanithy N, Black S: Patency rates after stenting across the inguinal ligament for treatment of post-thrombotic syndrome using nitinol venous stents. *J Vasc Surg Venous Lymphat Disord* 5(1):148, 2017
6. Lowe HJ, Ferris TA, Hernandez PM, Weber SC: STRIDE- an integrated standards-based translational research informatics platform. *AMIA Annu Symp Proc* 2009:391–395, 2009
7. Harris PA, Taylor R, Thielke R, Payne J, Gonzalez N, Conde JG: Research electronic data capture (REDCap) – A metadata-driven methodology and workflow process for providing translational research informatics support. *J Biomed Inform* 42(2):377–381, 2009
8. Cancer Institute Research Database (SCIRDB) Overview. Available at <https://med.stanford.edu/dbds/cool-tools/research-informatics-center/scirdb.html>. Accessed 5 April 2019.
9. Wickham H, Francios R, Henry L, Muller K.: Dplyr: A fast, consistent tool for working with data frame like objects, both in memory and out of memory. R package version 0.7.6. R Found Stat Comput, 2014.
10. Wickham H, Henry L: Tidy: An evolution of reshape2. Its designed specifically for data tidying (not general reshaping or aggregating) and works well with dplyr data pipelines. R package version 0.8.1. R Found Stat Comput, R Core Team, 2014.
11. R: A language and environment for statistical computing. R Found Stat. Comput.
12. Pre-processing GitHub Codebase. Available at https://github.com/LearningHealthcare/VITAL_Pre_Processing_Codebase. Accessed 21 August 2019.
13. Data Analysis GitHub Codebase. Available at https://github.com/LearningHealthcare/Inguinal_Ligament_Analysis_Codebase. Accessed 21 August 2019.

14. Therneau T, Lumley T: Survival: Survival analysis, including penalised likelihood. R package version 2.42–6 R found stat Comput, 2009.
15. Prinja S, Gupta N, Verma R: Censoring in clinical trials: Review of survival analysis techniques. *Indian J Community Med* 35(2): 217221, 2010
16. Mantel N: Evaluation of survival data and two new rank order statistics arising in its consideration. *Cancer Chemother Rep.* 50(3):16370, 1966
17. Wickham H, Chang W : ggplot2: Create elegant data Visualisations using the grammar of graphics. R package version 3.0.0 R found stat Comput, 2007.
18. RStudio Desktop. Available at rstudio.com/products/rstudio/#Desktop. Accessed 19 August 2019.
19. Vedantham S, Grassi CJ, Ferral H, Patel NH, Thorpe PE, Antonacci VP, Janne d'Othée BM, Hofmann LV, Cardella JF, Kundu S, Lewis CA, Schwartzberg MS, Min RJ, Sacks D: Reporting standards for endovascular treatment of lower extremity deep vein thrombosis. *J Vasc Interv Radiol* 20(7):S391S408, 2009

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.