

Interactive Exploration of Longitudinal Cancer Patient Histories Extracted From Clinical Text

Zhou Yuan, MS¹; Sean Finan, BS²; Jeremy Warner, MD, MS³; Guergana Savova, PhD²; and Harry Hochheiser, PhD¹

PURPOSE Retrospective cancer research requires identification of patients matching both categorical and temporal inclusion criteria, often on the basis of factors exclusively available in clinical notes. Although natural language processing approaches for inferring higher-level concepts have shown promise for bringing structure to clinical texts, interpreting results is often challenging, involving the need to move between abstracted representations and constituent text elements. Our goal was to build interactive visual tools to support the process of interpreting rich representations of histories of patients with cancer.

METHODS Qualitative inquiry into user tasks and goals, a structured data model, and an innovative natural language processing pipeline were used to guide design.

RESULTS The resulting information visualization tool provides cohort- and patient-level views with linked interactions between components.

CONCLUSION Interactive tools hold promise for facilitating the interpretation of patient summaries and identification of cohorts for retrospective research.

JCO Clin Cancer Inform 4:412-420. © 2020 by American Society of Clinical Oncology

Licensed under the Creative Commons Attribution 4.0 License 

INTRODUCTION

The complexities of cancer care create significant challenges for the extraction of information for retrospective research. As patients progress through diagnosis to treatment and subsequent monitoring, multiple encounters with varying specialists generate a rich set of clinical notes. For patients undergoing lengthy or multimodal (eg, a combination of surgery, chemotherapy, and radiotherapy) treatment, hundreds or thousands of notes can be generated along the “cancer journey.” Review of these notes can be a laborious interpretive challenge, often involving many hours of time for medical professionals who must read through collections of notes to prepare summarized abstractions in spreadsheets or databases. This process is also brittle, as reviews conducted for one study may miss items of potential interest to subsequent studies. Although ad hoc solutions such as the “oncologic history” have spontaneously developed as information collection devices, they are not necessarily universal, accurate, or complete.¹

The Cancer Deep Phenotype Extraction (DeepPhe) project is developing informatics solutions to overcome these inefficiencies. Unlike prior work applying

natural language processing (NLP) techniques to individual cancer documents,²⁻⁵ DeepPhe combines details from multiple documents to form longitudinal summaries. Classic and state-of-the-art NLP techniques for extracting individual concepts are used alongside a rich information model⁶ and techniques for care episode classification,⁷ cross-document coresolution,⁸ and rule-based inference to summarize diagnoses, treatments, responses, and temporal relationships as needed to support retrospective research.⁹ We expect that DeepPhe will be used either by clinicians or researchers with appropriate permissions to read notes de-identified by honest brokers or through other appropriate means. DeepPhe v3 was released in March 2019 and is available on GitHub.¹⁰

The application of the NLP tools to notes collected over months or years can lead to hundreds of observations: one modestly sized test data set of 49 patients had an average of > 245 facts/patient (standard deviation [SD], 99.3), spread over an average of 30.6 notes (SD, 18.4). Information visualization tools have the potential to help users easily interpret these rich records. The DeepPhe multilevel information model can easily support the “overview first, zoom and filter, details on demand”¹¹ approach that has proven successful in many previous visualization efforts. In the case of

Author affiliations and support information (if applicable) appear at the end of this article.

Accepted on April 7, 2020 and published at ascopubs.org/journal/cci on May 8, 2020; DOI <https://doi.org/10.1200/CCI.19.00115>

CONTEXT

Key Objective

How can interactive tools help researchers and clinicians understand longitudinal histories of patients with cancer extracted from notes via natural language processing?

Knowledge Generated

Web-based views at both the patient and cohort levels can provide display summary and detail information about complex cancer cases. Interaction techniques linking views at different granularities can enable navigation between summaries and details.

Relevance

Interactive tools for exploring summary representations of cases provide the possibility of easing interpretation of complex details as needed to inform care or to drive translational research.

DeepPhe, “details on demand” suggests drilling down from summarized representations to inference rules and specific spans of text that provide the provenance for those higher-level summaries.

Our goal is to develop a multiscale visualization to help researchers interpret the complexities of relationships between cancers, tumors, treatments, responses, biomarkers, and other key attributes. We draw on a substantial body of prior work on visual cohort extraction tools, many of which have used temporal or flow metaphors to characterize temporal trends or transitions across patient populations.¹²⁻¹⁵ The DeepPhe-Viz tool will extend these efforts with facilities for addressing challenges associated with the ambiguities of interpreting natural language. Our design of this tool was motivated by insights from qualitative inquiries with potential users and informed by our multilevel information model.

METHODS

Qualitative Inquiry

We conducted unstructured qualitative interviews with clinical cancer researchers at the University of Pittsburgh and Magee-Women’s Research Institute. Participants were a convenience sample identified through professional contacts of the research team. Interviews were conducted one on one, in participants’ workspaces, and covered a variety of questions focusing on challenges in cancer retrospective research, including goals, information needs, representations, bottlenecks, and challenges. Although contextual inquiry¹⁶ observations of researchers’ work as they reviewed clinical notes would have been preferred, interviewers did not have institutional review board clearance to see the de-identified patient data used by the researchers. Instead, interviews focused on general descriptions of the work and related challenges, including discussions of database schemas and tools such as spreadsheets used to manage extracted information.

All interviews were audio-recorded. Interviews were conducted and analyzed by a coauthor with extensive

experience in human-computer interaction research (H.H.), using an emergent coding approach¹⁷ to extract information needs, problems, design suggestions, and other relevant information. Comments were specifically reviewed to identify user challenges, classified through emergent code into those involving information availability, access, quality, and interpretation. Results from these analyses were used to develop user personae describing potential DeepPhe users, user stories involving specific tasks, competency questions detailing specific information requirements, and flow diagrams describing user processes. The University of Pittsburgh Human Research Protection Office classified these inquiries as exempt (PRO13120154).

Information Model

Qualitative inquiry results were used to develop an information model capable of representing relevant items and attributes at multiple granularities, ranging from individual text mentions to patient summaries.⁶

Mentions. Text spans in source documents covering concepts of interest, including tumors, body locations, treatments, stage indicators, biomarkers, and other key elements. Mentions have individual properties, such as negation, uncertainty, and historicity.

Compositions. Aggregations of mentions pertaining to the same unique entity or event. Composition Relations formulate clinical attributes and interconnection.

Episodes. Collection of documents in key event intervals, initially including work-up, diagnosis, medical decision making, treatment, and follow-up.

Patient Summary. Descriptions of cancers, tumors, treatments, and genomics, abstracted across the entire span of the patient history.

Natural Language Processing

Apache cTAKES¹⁸ pipelines were extended to extract individual mentions of cancer information, populating the mention level of the model. Mentions from each document

were aggregated and simplified via coreference resolution to form the composition level. Machine-learning models trained on annotated data are used to assign documents to episodes. Composition-level mentions are processed by a series of summarization rules to generate the high-level phenotypes. Results are stored in a Neo4j graph database.¹⁹ The initial DeepPhe architecture is described in detail by Savova et al.⁹

Visualization

Insights from qualitative inquiries informed the software requirement specification along with a corresponding series of low-fidelity prototypes for the interactive tools and visualizations. Subsequent iterative improvements of the functional software were informed by feedback from translational cancer researchers, cancer registrars not directly involved in the DeepPhe project, and oncologists, including coauthor J.W. Revisions focused on enhancing the multi-scale visualization capabilities (linking high-level summaries to individual text mentions)¹¹ and improving the interactive coordination between various views.²⁰

The DeepPhe-Viz tool was developed as a web application, using the Node.JS web platform,²¹ to provide a middle-ware layer capable of retrieving data through the Neo4j bolt protocol.²² The visualization interface was implemented in HTML, CSS, Javascript, and the D3 visualization toolkit.²³ The DeepPhe-Viz tool is available on GitHub.²⁴

RESULTS

Participants

Five researchers participated in the qualitative inquiries. Four had medical degrees, including two postdoctoral trainees, one practicing oncologist, and one full-time researcher. The fifth was a cancer epidemiologist with a PhD. Participants focused on either breast (n = 1) or ovarian (n = 4) cancer. Interviews were approximately 1 hour long.

Qualitative Inquiry and Visualization Requirements

User challenges identified during interviews involved difficulties with information availability, access, quality, and interpretation. Although some issues were specific to the types of cancer or the context of care, most were more broadly applicable (Table 1).

Together with informant descriptions of information needs and goals, these challenges informed the creation of user stories detailing specific tasks to be conducted for individual patients and/or at the cohort level. These user stories were broadly grouped into 14 requirement categories (Table 2).

Interactive Visualization Environment

As development of the DeepPhe NLP tools is an ongoing effort, prototype implementation of the visualization tools has been facilitated by the construction of synthetic details to complete fields that cannot yet be extracted by DeepPhe. The current prototype displays extracted results for cancer stage, diagnosis, treatments, tumor size, histologic type, tumor extent, cancer cell line, body site, and biomarkers. Synthesized results for date of birth and menopausal status are also displayed.

Cohort View

The DeepPhe cohort viewer (Fig 1) provides multiple complementary views. Tumor stages are shown in two views: a simple histogram of diagnosis stages (Supporting R7; Fig 1A) and an age distribution box plot (Fig 1B). A list of patient names (Fig 1C) enables quick identification of specific patients, and a scrollable list of diagnoses associated with each patient (Fig 1D) facilitates comparison between patients. Biomarker views (Fig 1E and 1F) show which patients have identified biomarkers and the distribution of observations among the active members of the cohort. The stage histogram can also be used to focus on individuals with specific stages: clicking on one of the bars will update the histogram and all other components to show only those items matching the stated criteria (Filter, R10). The double-thumb slider on the patient age by stage view (Fig 1B) also acts as a filter (R10). Each of these views also provides an overview of the associated distributions (R1).

Patient View

The DeepPhe-Viz patient view provides several panes at varying levels of granularity.

Under *patient details, cancer, and tumor*, overviews (R1) of patient demographics (Fig 2A) and cancer and tumor diagnoses (Fig 2B and 2C) are shown at the top left, providing a concise summary of patient details. Cancer details

TABLE 1. User Information Challenges Identified During Contextual Interviews

Category	Specific Challenges
Availability	Missing data for out-of-network patients; missing dates; progression/response to therapy missing; missing data regarding key questions or procedures (smoking, use of contraceptives, etc); limited follow-up information; unreported information (endometriosis in pathology reports); difficulty in finding healthy controls
Access	Information spread across multiple systems
Quality	Inconsistencies in drug name presentation challenge interpretation; abbreviations for conjunctions or alternative drugs
Interpretation	Reconciliation across report types; information about medications found in both physician and nursing reports; tumor registry data might not match records; contradictory information; differences in interpretation across groups; verification; extraction of multiple values for trend detection; interpretation of change in patients in the context of changes in practice (advent of neoadjuvant therapy); interpretation of outcomes of procedures.

TABLE 2. System Requirements and Associated User Stories

System Requirement	Cohort	Patient	User Stories	Implemented?
R1: Overview	✓	✓	Explore distribution of available data	✓
R2: Temporal	✓	✓	Note time points or intervals corresponding to changes in clinical practice or other factors that might influence results; examine trends; review care episodes	✓
R3: Text		✓	Refer to original clinical text	✓
R4: Provenance		✓	Identify original text used to extract/infer observations	✓
R5: Multirecords		✓	Use multiple types of records to bolster interpretations	✓
R6: Quality Assurance		✓	Verify accuracy and relevance of extracted data; check for consistency within document sets and between documents and unstructured or registry data	✓
R7: Stage	✓	✓	Identify cancer stage	✓
R8: Biomarkers & Genomics	✓	✓	Determine biomarker status and method used to make determination	✓
R9: Treatments		✓	Identify treatments provided; relate start and end dates of treatments to outcomes; understand response to treatment/progression: failure to respond, recurrence/no benefit, metastasis, developed resistance	Partial
R10: Filters	✓	✓	Filter cohorts based on categorical or temporal values; focus exploration based on relevant constraints	✓
R11: Uncertainty	✓	✓	Communicate both inherent ambiguity in notes and confidence in extracted information	
R12: Search	✓		Conduct searches based on clinical criteria and temporal relationships	
R13: Comparison	✓		Compare outcomes across selected cohorts	
R14: Criteria	✓		Report criteria used to identify or compare cohorts	

NOTE. Requirements not marked as implemented will be addressed in future releases.

include summary cancer attributes, cancer stage (R7: Stage), treatments (R9: Treatments), cell line, and TNM values.²⁵ The patient shown here has two independent cancer diagnoses, each providing overall stage summaries, treatments, and laterality, along with tumor details including specific diagnoses, biomarkers (R8: Biomarkers & Genomics), and other details shown as expandable lists of attributes colored to indicate classes of information. This approach provides a compact summary. A toggle at the top of the tumor summary pane supports switching to tabular views when desired. Tumor and cancer details can be selected to reveal individual text spans contributing to the summary element (R3: Text; R4: Provenance), thus providing an example of the use of the hierarchical model to go from summary to individual observation. Examination of these details can also be used to support R6: Quality Assurance.

The *clinical note timeline* supports R2: Temporal by displaying multiple types of notes (R5: Multirecords) on a timeline with one lane for each type of note (progress, radiology, and surgical pathology; Fig 2D). Notes are color-coded according to episode. A double-thumb scroll bar below the timeline allows zooming and panning across the extent, which spans from the interval between the first and last available documents. Episode labels above the timeline can be clicked to zoom the timeline to documents contained in the specified episode.

Below the timeline, the *explanation panel* (Fig 2E) supports R3: Text, R4: Provenance, and R6: QA by bridging the gap between the inferred attributes of the cancer and tumor summaries (Figs 2B and 2C) and the text of the clinical note (Fig 2G). Selection of summary items from the cancer or tumor summary lists leads to a display in the explanation panel describing the selected fact, along with information about its derivation from the given document. Like the tumor summary views, this panel illustrates the utility of the multilevel information model for moving between summary and individual observation, helping the user verify that the summarized assertion is indeed correct.

The *mention pane* (Fig 2F) provides a summary of mentions extracted from the selected document, supporting R3: Text and R4: Provenance. Each mention can be clicked to highlight the appropriate scan in the note view (Fig 2G), thus providing the user with additional tools for verifying correctness of the NLP output.

Navigation through multiple levels of abstraction is illustrated in Figure 2. The selection of tumor summary item “Ductal Breast Carcinoma in situ” (Fig 2B) led to the display of the “Invasive Ductal Carcinoma” in the explanation pane (Fig 2E) and the display of relevant mentions from Report 48 (Fig 2E). Clicking on the “positive” mention leads to text confirming the mention of invasive ductal carcinoma (Fig 2G).

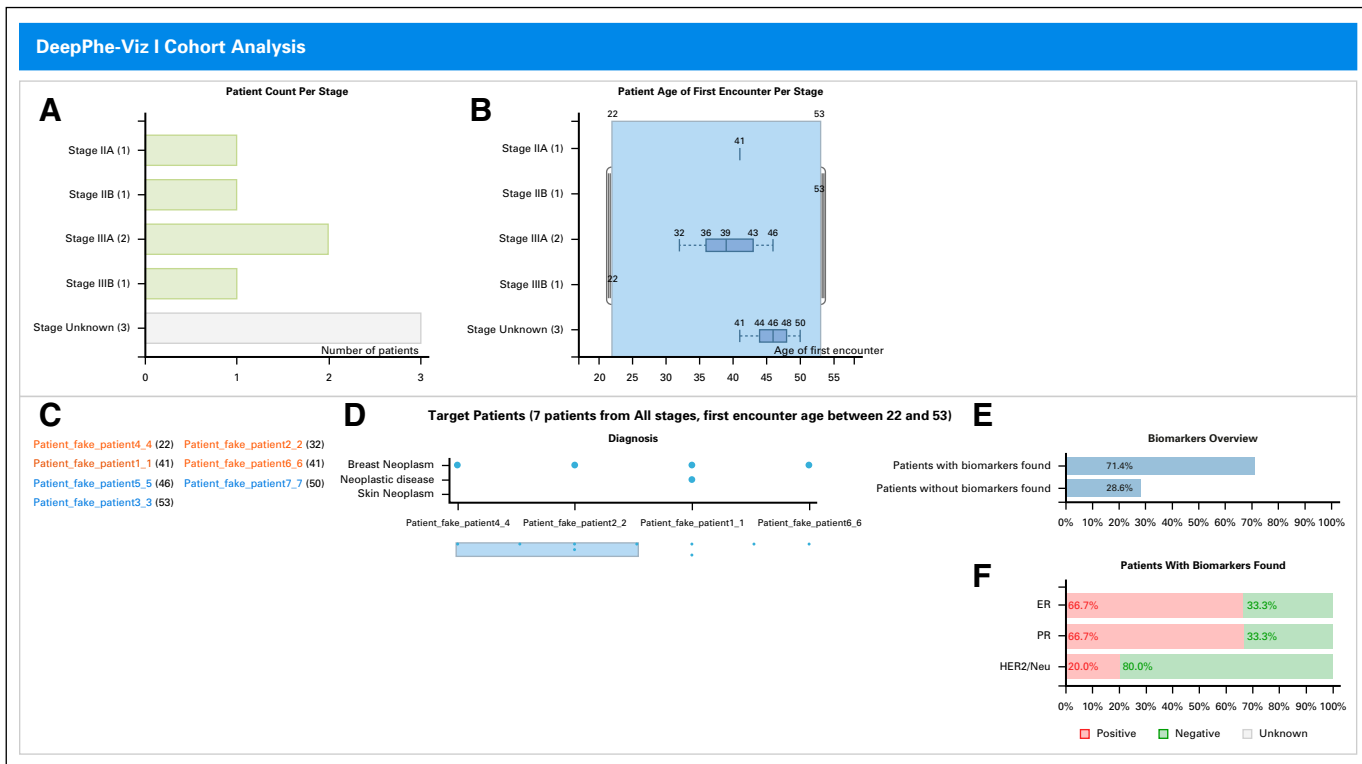


FIG 1. Cohort view: (A) Distributions of cancer stages, (B) distribution of ages, (C) scrollable list of patients, (D) distribution of diagnoses across patients, (E) distribution of biomarker counts, (F) status of three key breast cancer markers. As inference rules for summary staging are not included in DeepPhe (due to American Joint Committee on Cancer licensing requirements), a large portion of records are stage unknown.

DISCUSSION

The substantial amounts of clinical text associated with histories of patients with cancer present significant challenges for retrospective research. With histories involving dozens of relevant notes, manual expert review will not be sufficient for the large-scale analyses needed to drive innovation. Although advances in cross-document coreference²⁶ and other techniques currently being explored by the DeepPhe project show great promise in increasing the utility of clinical text, NLP is only a first step, providing an intermediate representation not directly consumable by end users. DeepPhe’s use of summarization and episode classification help provide order to the many facts that might be extracted from a set of patient records, but additional support is needed to turn these details into actionable understanding.

Our visualization tool is designed to tackle the four primary challenges associated with interpretation of these data: comparing patients (in the cohort view), facilitating exploration of patient histories over the time course of the available records, linking higher-level summaries to individual observations, and verifying output. Patient comparisons are necessary to enable identification of cohorts matching desired criteria. Aggregation of individual observations into higher-level clinically meaningful constructs will be necessary to easily answer key research questions

such as “which patients were treated with neoadjuvant therapy?” while linkages between those aggregations and individual text mentions enable verification of results, thus building user confidence in output.

The DeepPhe visualization tool represents a first step toward these goals, providing preliminary patient and cohort views of data for patients with cancer at multiple granularities. Although limited to a subset of desired data types, the current version illustrates basic functionality needed to address key requirements (Table 2) and outstanding challenges that have been identified during the evolution of the tools. Further engagement with domain experts representing multiple classes of stakeholders will be needed to ensure alignment between user needs and system functionality,

Unlike many previous text analytics tools that focus on classification²⁷ or more exploratory analysis of large text corpora,²⁸ the DeepPhe tools combine NLP results with an analytics interface, thus forming a complete analytics platform. DeepPhe is perhaps most similar to HARVEST,²⁹ which presents observations extracted from NLP in a timeline view. However, DeepPhe’s information model and inference rules provide support for cancer-specific higher-level abstractions not found in HARVEST. Future enhancements might include interactive features explored in related projects, including support for interactive

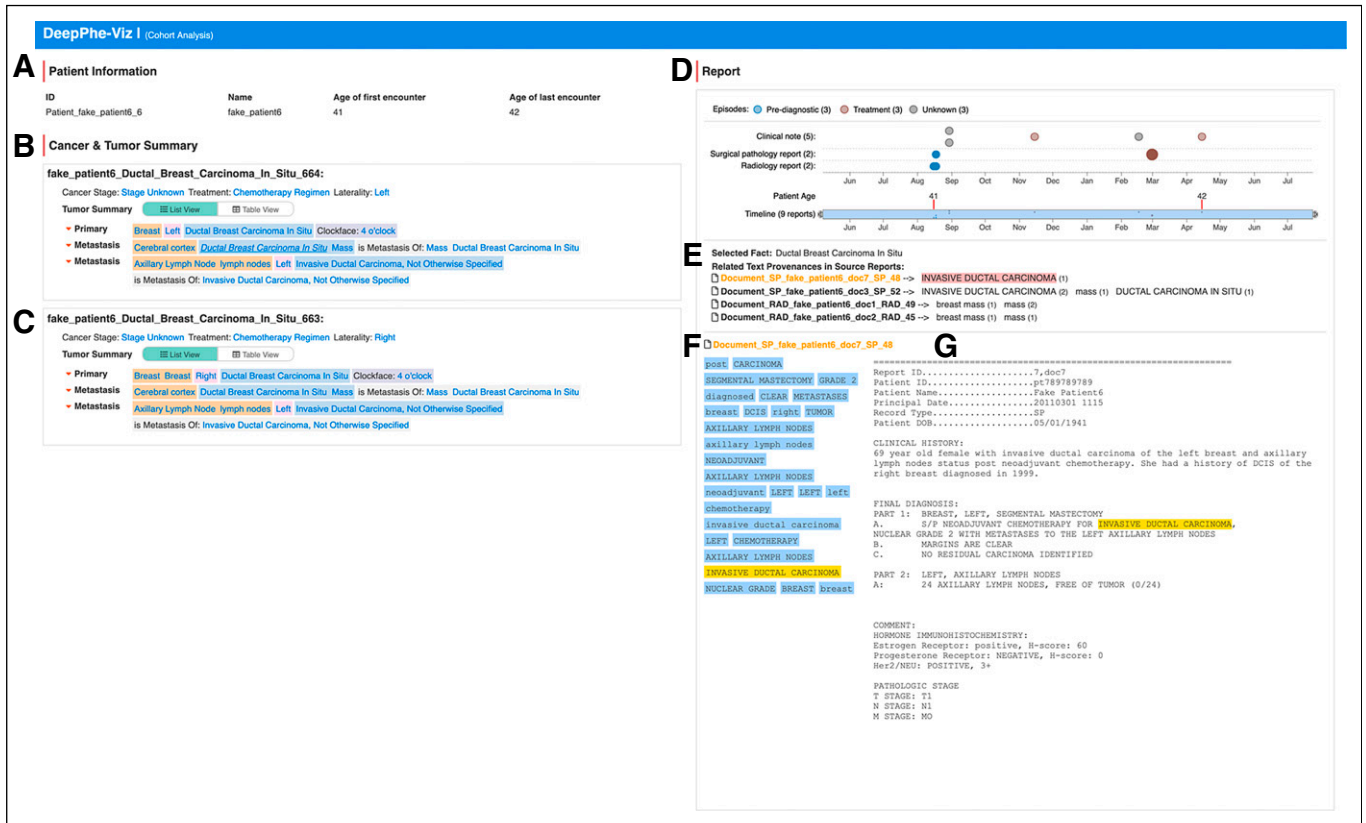


FIG 2. The patient visualization view, displaying manually fabricated synthetic data. (A) Patient demographics and cancer summary. (B, C) Cancer and Tumor summary for two different cancers. (D) Interactive timeline displaying documents by type and episode. (E) Inference provenance display. (F) Document-level mentions. (G) Clinical note. Selection of “Ductal Breast Carcinoma In Situ” in the cancer summary (B) leads to highlight of the first relevant note in the timeline. The provenance of this observation can be seen (E) in the display of relevant provenance rules, and (F) highlight of the selected term, (G) in the relevant document.

revisions of the NLP models,^{30,31} application to federated data sets,³² and additional visualizations display cohort-level patterns.¹²⁻¹⁵

Expanding the utility of the clinical text for identifying both cohorts and individual patients may aid in the interpretive process. Improved displays for both rendering and interpreting inference rules linking higher-order abstractions to individual text mentions may be helpful for complex inferences, particularly when cross-document inference is involved. Techniques for linking observations across documents will also prove useful for identifying recurring concepts identified through cross-document coreference resolution. At the cohort level, visualization of text patterns, perhaps enhanced through a Word Tree³³ or similar visualization, might help users interpret key phrases indicative of observations of interest.

DeepPhe visualization functionality will evolve alongside NLP capabilities. Although extraction and classification of individual mentions has led to promising results in many of the attributes currently shown in the prototype visualizations, much work remains to be done in the inference of higher-level aggregations and, subsequently, the inclusion

of these representations in the visualization. Two key examples involve tumors and treatments. Linking multiple tumor references across temporal extents, and including these intervals in the timeline view, will provide valuable perspective on cancer progression and response.

Enhanced temporal aggregation will also drive extensions of the DeepPhe cohort view. Incorporation of per-document episode enhancement techniques, alongside orderings of treatments and time spans of specific tumors, will support temporally aligned cohort analysis using techniques similar to those used in Outflow,¹² Frequence,¹⁴ EventFlow,¹³ and related systems.¹⁵ Temporal¹³ and logical³⁴ search facilities are also planned, with pattern search^{15,35} a possibility for future work. Similar to previous tools focused on specific domains³⁶ or care pathways and treatment plans,³⁷⁻³⁹ we will use episode annotations and the semantics of the DeepPhe information model to focus designs on the specific challenges of interpreting cancer data.

Inclusion of treatment information, particularly for chemotherapeutic regimens, may provide investigators with insights into treatment histories and possible impacts. Effectively displaying treatments will require inference not

only of specific start and stop times of various drugs but ideally of identification or inference of multidrug protocols. Extension of DeepPhe NLP tools to identify medication regimens on the basis of the HemOnc ontology^{40,41} is a high priority.

As DeepPhe interactive tools evolve to include these new data elements, appropriate handling of uncertainty and missing information will become increasingly critical. NLP temporal modeling techniques^{8,42} might be used in combination with structured electronic health record data to eliminate some ambiguity, but many details will likely remain unspecified. Cohort and patient tools will need both appropriate display of these underspecified constraints and appropriate semantics for any related queries or filters. Temporal ambiguities also underscore the importance of tools for explicitly describing search criteria and for facilitating comparisons between cohorts as techniques that might reduce the risk of misinterpretation.

Evaluation of analytic tools such as the DeepPhe visualizations has been the subject of an active body of research. As information visualization tasks are often exploratory and ill defined, traditional metrics such as task completion time and accuracy may not be particularly informative, leading to the need for investigations into descriptions of the use of the tool in terms of analytic processes used,

types of interactions, and clarity of explanations of data.⁴³ Insight-based evaluations aimed at quantifying novel understanding might also be considered.⁴⁴ Planned evaluations will follow a phased approach, combining small-scale usability visualizations with larger laboratory studies and eventual observations of the use of the tool in context.⁴⁵ DeepPhe and the DeepPhe-Viz tool are available on GitHub.^{10,24}

DeepPhe's current architecture assumes that text files to be processed are available in a simple file structure, organized by patient. Efforts to provide DeepPhe functionality through application programming interfaces, allowing integration with other tools and data warehouse environments, are underway.

Facilitating the interpretation of complex, longitudinal patient histories is an important challenge for understanding cancer treatment and outcomes. The DeepPhe project uses a multifaceted approach, combining NLP, inference, information modeling, and interactive visualizations to provide researchers with detailed descriptions that span the gap between key phenomena of interest and specific documentary evidence. Extension of proposed prototype designs to handle richer data, particularly involving temporal spans, will set the stage for deployment with clinical researchers and subsequent evaluation studies.

AFFILIATIONS

¹University of Pittsburgh, Pittsburgh, PA

²Boston Children's Hospital, Boston, MA

³Vanderbilt University, Nashville, TN

Preprint version available on [bioRxiv](https://www.biorxiv.org/).

CORRESPONDING AUTHOR

Harry Hochheiser, PhD, 5607 Baum Blvd, Pittsburgh, PA 15206; e-mail: harryh@pitt.edu.

PRIOR PRESENTATION

Presented at the 2018 Workshop on Visual Analytics in Health Care, San Francisco, CA, November 3, 2018.

SUPPORT

Supported by National Cancer Institute Grant No. 1U24CA184407.

AUTHOR CONTRIBUTIONS

Conception and design: Zhou Yuan, Guergana Savova, Harry Hochheiser

Financial support: Harry Hochheiser

Administrative support: Harry Hochheiser

Provision of study material or patients: Harry Hochheiser

Collection and assembly of data: Zhou Yuan, Sean Finan, Harry Hochheiser

Data analysis and interpretation: Sean Finan, Jeremy Warner, Harry Hochheiser

Manuscript writing: All authors

Final approval of manuscript: All authors

Accountable for all aspects of the work: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated unless otherwise noted. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/cci/author-center.

Open Payments is a public database containing information reported by companies about payments made to US-licensed physicians ([Open Payments](https://www.fda.gov/oc/openpayments/)).

Jeremy Warner

Stock and Other Ownership Interests: HemOnc.org

Consulting or Advisory Role: Westat, IBM

Travel, Accommodations, Expenses: IBM

No other potential conflicts of interest were reported.

ACKNOWLEDGMENT

Rebecca Jacobson, Eugene Tseytlin, Girish Chavan, Olga Medvedeva, Kevin Mitchell, and Melissa Castine made important contributions to early phases of this work. Alicia Beeghly-Fadiel, Eric Durbin, Elizabeth Buchbinder, and others provided valuable input into the design of the tool.

REFERENCES

1. Warner JL, Anick P, Hong P, et al: Natural language processing and the oncologic history: Is there a match? *J Oncol Pract* 7:e15-e19, 2011
2. AAIAbdulsalam AK, Garvin JH, Redd A, et al: Automated extraction and classification of cancer stage mentions from unstructured text fields in a central cancer registry. *AMIA Jt Summits Transl Sci Proc* 2017:16-25, 2018
3. Warner JL, Levy MA, Neuss MN, et al: ReCAP: Feasibility and accuracy of extracting cancer stage information from narrative electronic health record data. *J Oncol Pract* 12:157-158, e169-e7, 2016
4. Xie F, Lee J, Munoz-Plaza CE, et al: Application of text information extraction system for real-time cancer case identification in an integrated healthcare organization. *J Pathol Inform* 8:48, 2017
5. Kreimeyer K, Foster M, Pandey A, et al: Natural language processing systems for capturing and standardizing unstructured clinical information: A systematic review. *J Biomed Inform* 73:14-29, 2017
6. Hochheiser H, Castine M, Harris D, et al: An information model for computable cancer phenotypes. *BMC Med Inform Decis Mak* 16:121, 2016
7. Lin C, Miller T, Amiri H, et al: Classification of electronic medical records of breast cancer and melanoma patients into clinical episodes. Presented at the 30th Anniversary AACR Special Conference Convergence: Artificial Intelligence, Big Data, and Prediction in Cancer, Newport, RI, October 14-17, 2018
8. Lin C, Miller T, Dligach D, et al: A BERT-based universal model for both within- and cross-sentence clinical temporal relation extraction. Proceedings of the 2nd Clinical Natural Language Processing Workshop. Stroudsburg, PA, Association for Computational Linguistics, 2019, pp 65-71
9. Savova GK, Tseytlin E, Finan S, et al: DeepPhe: A natural language processing system for extracting cancer phenotypes from clinical records. *Cancer Res* 77:e115-e118, 2017
10. GitHub: DeepPhe-Release. <https://github.com/DeepPhe/DeepPhe-Release>
11. Shneiderman B: The eyes have it: A task by data type taxonomy for information visualizations. Proceedings of the 1996 IEEE Symposium on Visual Languages. Washington, DC, IEEE Computer Society, 1996, pp 336
12. Wongsuphasawat K, Gotz D: Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Trans Vis Comput Graph* 18:2659-2668, 2012
13. Monroe M, Lan R, Lee H, et al: Temporal event sequence simplification. *IEEE Trans Vis Comput Graph* 19:2227-2236, 2013
14. Perer A, Wang F: Frequence: Interactive mining and visualization of temporal frequent event sequences. Proceedings of the 19th International Conference on Intelligent User Interfaces. New York, NY, Association for Computing Machinery, 2014, pp 153-162
15. Gotz D, Wang F, Perer A: A methodology for interactive mining and visual analysis of clinical event patterns using electronic health record data. *J Biomed Inform* 48:148-159, 2014
16. Beyer H, Holtzblatt K: Contextual Design: Defining Customer-Centered Systems. San Francisco, CA, Morgan Kaufmann, 1998
17. Lazar J, Feng JH, Hochheiser H: Research Methods in Human-Computer Interaction. San Francisco, CA, Morgan Kaufmann, 2017
18. Savova GK, Masanz JJ, Ogren PV, et al: Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): Architecture, component evaluation and applications. *J Am Med Inform Assoc* 17:507-513, 2010
19. Neo4j: Neo4j graph platform. www.neo4j.com
20. Roberts JC: State of the art: Coordinated multiple views in exploratory visualization. Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007). Zurich, Switzerland, IEEE, 2007, pp 61-71
21. OpenJS Foundation: Node.js. www.nodejs.org
22. Bolt protocol. <https://boltprotocol.org/>
23. Bostock M, Ogievetsky V, Heer J: D³: Data-driven documents. *IEEE Trans Vis Comput Graph* 17:2301-2309, 2011
24. GitHub: DeepPhe-Viz. <https://github.com/deepphe/deepphe-viz>
25. Gospodarowicz MK, Brierley JD, Wittekind C: TNM Classification of Malignant Tumours. Oxford, United Kingdom, John Wiley & Sons, 2017
26. Miller T, Dligach D, Bethard S, et al: Towards generalizable entity-centric clinical coreference resolution. *J Biomed Inform* 69:251-258, 2017
27. Heimerl F, Koch S, Bosch H, et al: Visual classifier training for text document retrieval. *IEEE Trans Vis Comput Graph* 18:2839-2848, 2012
28. Görg C, Liu Z, Stasko J: Reflections on the evolution of the Jigsaw visual analytics system. *Inf Vis* 13:336-345, 2014
29. Hirsch JS, Tanenbaum JS, Lipsky Gorman S, et al: HARVEST, a longitudinal patient record summarizer. *J Am Med Inform Assoc* 22:263-274, 2015
30. Trivedi G, Pham P, Chapman WW, et al: NLPReViz: An interactive tool for natural language processing on clinical text. *J Am Med Inform Assoc* 25:81-87, 2018
31. Trivedi G, Dadashzadeh ER, Handzel RM, et al: Interactive NLP in clinical care: Identifying incidental findings in radiology reports. *Appl Clin Inform* 10:655-669, 2019
32. Crowley RS, Castine M, Mitchell K, et al: caTIES: A grid based system for coding and retrieval of surgical pathology reports and tissue specimens in support of translational research. *J Am Med Inform Assoc* 17:253-264, 2010
33. Wattenberg M, Viégas FB: The Word Tree, an interactive visual concordance. *IEEE Trans Vis Comput Graph* 14:1221-1228, 2008
34. Krause J, Perer A, Stavropoulos H: Supporting iterative cohort construction with visual temporal queries. *IEEE Trans Vis Comput Graph* 22:91-100, 2016
35. Shknevsky A, Shahar Y, Moskovitch R: Consistent discovery of frequent interval-based temporal patterns in chronic patients' data. *J Biomed Inform* 75:83-95, 2017
36. Huang C-W, Syed-Abdul S, Jian W-S, et al: A novel tool for visualizing chronic kidney disease associated polymorbidity: A 13-year cohort study in Taiwan. *J Am Med Inform Assoc* 22:290-298, 2015
37. Aigner W, Miksch S: CareVis: Integrated visualization of computerized protocols and temporal patient data. *Artif Intell Med* 37:203-218, 2006
38. Bettencourt-Silva JH, Mannu GS, de la Iglesia B: Visualisation of integrated patient-centric data as pathways: Enhancing electronic medical records in clinical practice, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). Cham, Switzerland: Springer, 2016, pp 99-124
39. Gschwandtner T, Aigner W, Kaiser K, et al: CareCruiser: Exploring and visualizing plans, events, and effects interactively. Proceedings of the IEEE Pacific Visualization Symposium 2011 (PacificVis 2011). Los Alamitos, CA, IEEE Computer Society, 2011, pp 43-50
40. Mally AM, Jain SK, Yang PC, et al: Computerized approach to creating a systematic ontology of hematology/oncology regimens. *JCO Clin Cancer Inform* 10.1200/CCI.17.00142
41. Warner JL, Cowan AJ, Hall AC, et al: HemOnc.org: A collaborative online knowledge platform for oncology professionals. *J Oncol Pract* 11:e336-e350, 2015
42. Lin C, Dligach D, Miller TA, et al: Multilayered temporal modeling for the clinical domain. *J Am Med Inform Assoc* 23:387-395, 2016

43. Scholtz J, Plaisant C, Whiting M, et al: Evaluation of visual analytics environments: The road to the Visual Analytics Science and Technology Challenge evaluation methodology. *Inf Vis* 13:326-335, 2014
44. Saraiya P, North C, Duca K: An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Trans Vis Comput Graph* 11:443-456, 2005
45. Wu DTY, Chen AT, Manning JD, et al: Evaluating visual analytics for health informatics applications: A systematic review from the American Medical Informatics Association Visual Analytics Working Group Task Force on Evaluation. *J Am Med Inform Assoc* 26:314-323, 2019

