

SMART precision cancer medicine: a FHIR-based app to provide genomic information at the point of care

RECEIVED 31 August 2015
 REVISED 13 January 2016
 ACCEPTED 26 January 2016
 PUBLISHED ONLINE FIRST 27 March 2016



Jeremy L Warner,^{1,2} Matthew J Rieth,^{1,2} Kenneth D Mandl,^{3,4} Joshua C Mandel,^{3,4} David A Kreda,⁵ Isaac S Kohane,^{3,4,6,7} Daniel Carbone,⁸ Ross Oreto,⁸ Lucy Wang,⁸ Shilin Zhu,⁹ Heming Yao,¹⁰ Gil Alterovitz^{3,4,6}

ABSTRACT

Background Precision cancer medicine (PCM) will require ready access to genomic data within the clinical workflow and tools to assist clinical interpretation and enable decisions. Since most electronic health record (EHR) systems do not yet provide such functionality, we developed an EHR-agnostic, clinico-genomic mobile app to demonstrate several features that will be needed for point-of-care conversations.

Methods Our prototype, called Substitutable Medical Applications and Reusable Technology (SMART)[®] PCM, visualizes genomic information in real time, comparing a patient's diagnosis-specific somatic gene mutations detected by PCR-based hotspot testing to a population-level set of comparable data. The initial prototype works for patient specimens with 0 or 1 detected mutation. Genomics extensions were created for the Health Level Seven[®] Fast Healthcare Interoperability Resources (FHIR)[®] standard; otherwise, the prototype is a normal SMART on FHIR app.

Results The PCM prototype can rapidly present a visualization that compares a patient's somatic genomic alterations against a distribution built from more than 3000 patients, along with context-specific links to external knowledge bases. Initial evaluation by oncologists provided important feedback about the prototype's strengths and weaknesses. We added several requested enhancements and successfully demonstrated the app at the inaugural American Society of Clinical Oncology Interoperability Demonstration; we have also begun to expand visualization capabilities to include cancer specimens with multiple mutations.

Discussion PCM is open-source software for clinicians to present the individual patient within the population-level spectrum of cancer somatic mutations. The app can be implemented on any SMART on FHIR-enabled EHRs, and future versions of PCM should be able to evolve in parallel with external knowledge bases.

Keywords: health information management, electronic health records, genomics, neoplasms, information science, mobile health

BACKGROUND AND SIGNIFICANCE

The definition of cancer and the care of cancer patients are increasingly being driven by tumor genomics, aka molecular profiling.^{1,2} As the number of clinically relevant findings with prognostic implications rapidly expands, human cognitive capacity, as predicted in 1989,³ will no longer be able to keep up. Vanderbilt University Medical Center (VUMC) was an early adopter of near-universal genotyping of cancer specimens for a number of disease-specific “actionable” mutations, through the use of SNaPshot multiplexed PCR mutation panels.^{4–7} Even with the limited number of genes tested in SNaPshot, it is already very difficult for an individual clinician to be intimately familiar with the population distributions of genomic alterations and their implications.⁸ These difficulties will only grow worse as next-generation sequencing (NGS) will utilize cancer gene panels approaching 500 genes and potentially thousands of variants per specimen.^{9–11} Unfortunately, commercially available electronic health records do not provide ready means to display clinical genomic data, nor additional functionalities such as links to external knowledge bases including My Cancer Genome, a curated oncology gene variant knowledge base, among others. This gap must be addressed for clinicians to achieve the goals set forth in President Barack Obama's Precision Medicine Initiative.¹²

OBJECTIVE

We sought to develop an open-source application based on the Substitutable Medical Applications and Reusable Technology (SMART)

Health IT platform (www.smarthealthit.org), an open-access application programming interface (API) that enables apps to run broadly across the health care ecosystem.¹³ We used Health Level Seven International (HL7[®])'s Fast Healthcare Interoperability Resources (FHIR[®]) standard and its extension features for native representation of molecular profile data. The purpose of the resulting SMART Precision Cancer Medicine (PCM) app is to present population-level genomic health information to oncologists and their patients in real time as a component of clinical practice. We also wanted to demonstrate the ease of including seamless links to external knowledge bases within the app.

MATERIALS AND METHODS

The SMART API, SMART on FHIR, and genomics extensions

The SMART platform has been previously described.^{14,15} HL7 is a standards development organization that has developed several widely used standards in the health care space. The newest HL7 standard is FHIR, currently a draft standard for trial use. Similarly to HL7 V3, FHIR is a constraint on the HL7 Reference Information Model,¹⁶ with some minor modifications. In addition, FHIR is based around the latest Web technologies, such as representational state transfer APIs, and can be represented in extensible markup language, Java script object notation, or the resource description framework.

SMART on FHIR provides an app platform for health applications that integrates with EHRs, patient portals, personal health records,

Correspondence to Jeremy L. Warner, Assistant Professor of Medicine and Biomedical Informatics, Vanderbilt University, 2220 Pierce Ave, Preston Research Building 777, Nashville, TN 37232, USA; jeremy.warner@vanderbilt.edu; Tel: 1 (615) 322-5464; Fax: 1 (615) 343-7602 For numbered affiliations see end of article.

© The Author 2016. Published by Oxford University Press on behalf of the American Medical Informatics Association. All rights reserved. For Permissions, please email: journals.permissions@oup.com

and data warehouses. There are 3 key aspects of SMART on FHIR: (1) a data access layer based on FHIR, combined with a set of constraining profiles that lock down optionality and align vocabularies with Meaningful Use requirements,^{17,18} (2) a security layer that provides narrowly scoped authorization to specific portions of a patient's record via OAuth 2.0,¹⁹ and (3) a single-sign-on layer using OpenID Connect.²⁰ SMART on FHIR apps can be integrated into the context of an existing EHR or patient portal session, conveying the current patient, encounter, and other details of the host environment, or they can launch independently, such as on a mobile phone or device.

The SMART on FHIR Genomics API provides additional functionality to SMART on FHIR by extending the FHIR Observation resource to support clinical genomic data.²¹ These extensions are the basis of the FHIR Standard Profile for Genetics, which was published on September 23, 2015 (<http://www.hl7.org/FHIR/observation-genetics-cg-prf-1a.html>).

Translating local data into FHIR-compliant data

We were interested in displaying the following information for clinical consumption: (1) demographics including name, medical record number, gender, and age, (2) primary cancer diagnosis, and (3) molecular profile results. Comorbidities were not displayed at this pilot stage. Solid tumor oncology patients with certain histologies (eg, lung cancer, melanoma) seen at VUMC have, since July 2010, routinely undergone molecular profiling with SNaPshot, a fast, high-throughput, multiplex mutational profiling method based on the Applied Biosystems SNaPshot platform.^{5,7} SNaPshot assays for common somatic mutations across multiple cancer-associated genes, as defined in the Catalogue of Somatic Mutations in Cancer (COSMIC).²² SNaPshot panels have been performed on more than 4500 cancer specimens at VUMC (as of August 2015).

SNaPshot data are stored in the VUMC EHR and related data warehouses using an internally developed local code set. We transformed these data into unambiguous codable concepts after reviewing various terminology options for diseases, genes, gene alterations, and protein alterations (see [Supplementary Table 1](#)). This candidate list was generated by searching for codes for a representative disease (melanoma), gene (BRAF), and variant (BRAF p.V600E) in the National Cancer Institute Thesaurus, a comprehensive terminology that is used by the US Food and Drug Administration, the Clinical Data Interchange Standards Consortium, and genomic projects such as the Cancer Genome Atlas. We then reviewed all mappings through the

NCImetathesaurus to find additional terminologies. We also Web-crawled the external links on Gene Wiki, as well as links on the linked pages, to identify additional terminologies. Finally, we utilized a curated list created by HL7 membership (courtesy of Dr Clem McDonald). With this terminology identification complete, we arrived at a consensus through internal discussions as well as discussions with the HL7 Clinical Genomics work group as follows: diseases were mapped from ICD-9-CM to the Systematized Nomenclature of Medicine, Clinical Terms. Gene names were represented in Human Gene Naming Consortium nomenclature²³ and coded using the National Cancer Institute Thesaurus.²⁴ Gene reference sequence was represented in Consensus Coding Sequence terminology and coded using the National Center for Biotechnology Information's Consensus Coding Sequence database.²⁵ Gene alterations and predicted protein alterations were represented directly using Human Genome Variant Society compliant syntax.²⁶ For each gene, we also determined the relevant URL in the Gene Wiki knowledge base.²⁶ Examples of this mapping process are shown in [Table 1](#). Once mapping was complete, we created permanent FHIR instances for all patients who underwent SNaPshot testing by making extensions on the FHIR Observation and Specimen resources to support the required gene, reference sequence, and variant data, and stored these instances on a SMART on FHIR server. To keep our prototype requirements manageable, data from the minority of patients with multiple detected mutations were omitted. An example of the Java script object notation FHIR representation used to communicate with the app is shown in [Figure 1](#) and [Supplementary Table 2](#).

Display name was chosen to coordinate with clinicians' expectations of how genetic information is presented. Since SNaPshot is a DNA mutation assay, amino acid changes are predicted. BRAF: B-Raf proto-oncogene, serine/threonine kinase; EGFR: epidermal growth factor receptor; KRAS: Kirsten rat sarcoma viral oncogene homolog.

SMART PCM app overview

A team of software developers within the Vanderbilt-Ingram Cancer Center (VICC)'s Research Informatics Core developed a native iOS app optimized for the iPad and iPad Mini tablet devices (Apple Inc., Cupertino, CA, USA). Development proceeded in a continuous delivery build cycle with frequent input from clinical subject matter experts, especially with respect to how to properly model the clinical genomic data; the actual implementation of FHIR was straightforward. An integral component of the development process was implementation and

Table 1: Examples of local code mapping to standardized codable concepts.

Local Code	Human Gene Naming Consortium ID	National Cancer Institute Thesaurus	Consensus Coding Sequence	Human Genome Variant Society	Display Name	Gene Wiki URL (http://en.wikipedia.org/wiki/~)
BV600	HGNC:1097	C18363	5863.1	c.1799_1800TG>AA	BRAF V600E	~BRAF_(gene)
				p.V600E		
E790M	HGNC:3236	C17757	5514.1	c.2369C>T	EGFR T790M	~Epidermal_growth_factor_receptor
				p.T790M		
G12CM	HGNC:6407	C25785	8702.1	c.34G>T	KRAS G12C	~KRAS
				p.G12C		
G13CM	HGNC:6407	C25785	8702.1	c.37G>T	KRAS G13C	~KRAS
				p.G13C		

Figure 1: A snippet of the JSON FHIR code for a patient with lung cancer and a p.T790M mutation detected in the epidermal growth factor receptor gene. Three extensions to the FHIR Observation Resource are shown: (1) `assessed.gene`, which uses the NCI Thesaurus to represent the gene name in HGNC-compliant format; (2) `assessed.referenceSeq`, which uses the CCDS database to represent the gene reference sequence; (3) `assessed.variant`, which represents the observed gene mutation (c.2369C>T) and predicted protein alteration (p.T790M) directly in Human Genome Variant Society syntax. The full code for this patient is available in Supplementary Table 2.

```

100 -
101 -
102 -
103 -
104 -
105 -
106 -
107 -
108 -
109 -
110 -
111 -
112 -
113 -
114 -
115 -
116 -
117 -
118 -
119 -
120 -
121 -
122 -
123 -
124 -
125 -
126 -
127 -
128 -
129 -
130 -
131 -
132 -
133 -
134 -
135 -
136 -
137 -
{
  "url": "https://fhir/Mutation#assessed.gene",
  "valueCodeableConcept": {
    "coding": [
      {
        "system": "http://ncit.nci.nih.gov",
        "code": "C17757",
        "display": "EGFR"
      }
    ]
  }
},
{
  "url": "https://fhir/Mutation#assessed.referenceSeq",
  "valueCodeableConcept": {
    "coding": [
      {
        "system": "http://ncbi.nlm.nih.gov/CCDS",
        "code": "5514.1",
        "display": "CCDS 5514.1"
      }
    ]
  }
},
{
  "url": "https://fhir/Mutation#assessed.variant",
  "valueCodeableConcept": {
    "coding": [
      {
        "system": "http://www.hgvs.org/mutnomen",
        "code": "c.2369C>T",
        "display": "c.2369C>T"
      },
      {
        "system": "http://www.hgvs.org/mutnomen",
        "code": "p.T790M",
        "display": "T790M"
      }
    ]
  }
}
}

```

refinement of the app at the FHIR Connectathon 7 (September 13–14, 2014) and a Connectathon hosted by the American Society of Clinical Oncology (May 4–5, 2015); details are in [Supplementary Materials](#).

The PCM app is accessed from an icon on the home screen of an iPad. The user authenticates by using their standard VUMC username and password. As required by the OAuth 2.0 process, the user must authorize the app for initial data access. A splash-screen is presented and followed by a simple query interface, which allows a practitioner to look up a patient by name or medical record number.

A successful patient query brings up the patient-centric view, as shown in [Figure 2](#). In the example, the patient has lung cancer that harbors a KRAS p.G12C mutation. To prevent visual clutter in the initial series of pie charts, only common occurrences are labeled. The clinician can see the fully labeled information by interacting with the charts, as shown in [Figure 3](#). The “dive-in” detail includes full labels and exact patient counts. From that view, the user can link to external Gene Wiki content by touching the desired gene ([Figure 4](#)).

User feedback survey

After completion of the initial development cycle, which included FHIR Connectathon 7, we performed user testing with oncology clinicians. The SMART PCM app was evaluated in the VICC medical oncology clinic, with the data source being ~3800 patients who had undergone SNaPshot testing with 0 or 1 mutation detected. All queries for names,

medical record numbers, and genomic data were performed in real time from a local SMART on FHIR server, and data was fed to the app via FHIR bundle and resource objects. This initial evaluation focused on collecting clinician feedback and was not carried out in the presence of patients. Similar to other “apps” in the VUMC ecosystem (eg, WebPACS), the app is launched independently from the EHR and can be used synchronously or asynchronously per user needs. Two groups of users were targeted for evaluation of the application: fellows within the oncology training program and practicing attending oncologists. Fourteen users (9 fellows and 5 attendings) were approached for user evaluation, with a goal to ascertain the majority of problems and experiences for each group.²⁷ After they had freely explored the app’s functionality without a time limit, they were directed to complete a short online survey. Survey questions were created based on informal needs assessments and conversations between 2 clinical oncologists. Questions were created around user concepts and specific potential improvements. Additionally, users were able to enter open-ended free-text comments. All users completed the survey, which is described in the [Supplementary Materials](#).

Statistical Analysis

Survey responses were exported from Research Electronic Data Capture into the R statistical package for analysis. Free-text comments were coded by 2 independent raters (J.L.W., M.J.R.) for content, with

Figure 2: Example output of the SMART PCM app, showing a lung cancer patient with KRAS p.G12C mutation in the context of other lung cancer patients tested at VUMC. Further information is available to the user through interaction with the pie charts, all pieces of which are activated by touch. On the left, a pie chart shows the population distribution of gene mutations. In this example, it is evident that slightly more than half the patients have no mutation detected, whereas KRAS is the most commonly mutated gene. On the right, the distribution of variants of the mutated gene is shown, where it is evident that p.G12C is the most common KRAS mutation. In a case where a patient has no mutation detected, the variant pie chart is suppressed. Patient details (name, age, gender) are redacted to preserve PHI.



discrepancies adjudicated by joint consultation; interrater reliability was calculated using Cohen's kappa (κ). Hypothesis testing was performed using Fisher's exact test and Mann-Whitney U; all statistical tests were 2-sided.

Ethics and software availability

The described tool is a quality improvement initiative with the intent to implement a practice to improve the quality of patient care, and was thus determined by the Vanderbilt Institutional Review Board to be non-research per section 45 CFR 46.102(d) of the Health Insurance Portability and Accountability Act.

The source code for the app has been made openly available on GitHub (GitHub Inc., San Francisco, CA, USA): <https://github.com/dcarbon-smart-precision-cancer-medicine>. The source code for the branched SMART on FHIR server is also available: <https://github.com/ross-oreto/api-server>. Further information is available on the SMART App Gallery (<https://gallery.smarthealthit.org/vanderbilt-university-medical-center/smart-precision-cancer-medicine>).

RESULTS

Implementing SMART on FHIR Genomics at VUMC

VUMC's extensive research data warehouse is a relational database containing detailed information on over 3 million patients, dating back to 1992.²⁸ This includes most data from the EHR, such as clinical progress notes and provider-patient communications, as well as data feeds directly from laboratory and billing systems. Our app was built to interface with a local SMART on FHIR server, using data directly from the research data warehouse, analogous to previous work with i2b2.²⁹ The local server was similar to the SMART on FHIR prototype

described above, except that OAuth 2.0 was linked to a custom internal authentication and role-based authorization service, which utilizes lightweight directory access protocol for access to client applications. Access to the server was only enabled within the VUMC firewall and was audited.

User feedback

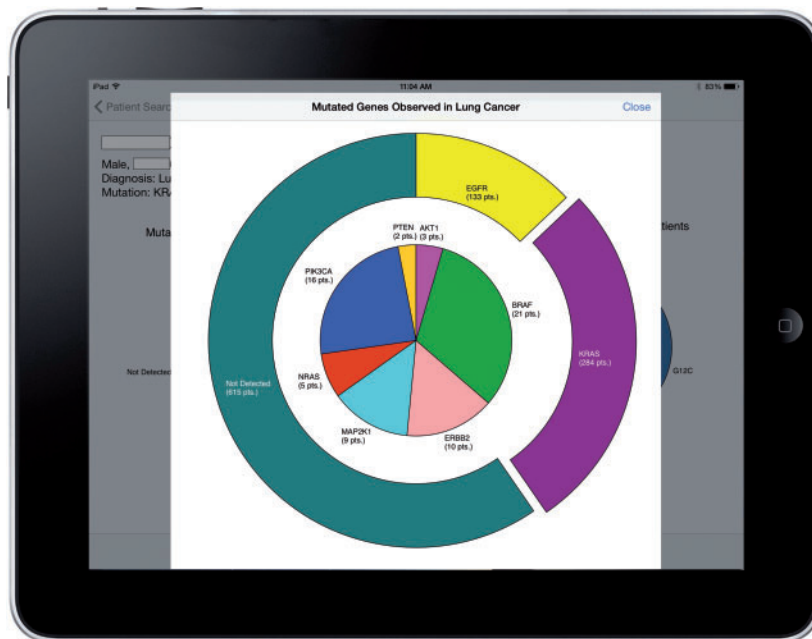
User evaluations were highly variable across all nominal domains tested, with no correlation among variables (eg, users who found the system quick and easy to use would not necessarily use it clinically; see [Supplementary Figure 1](#)). There was no statistical difference between fellows and attendings in their responses to any of the questions ($P > 0.05$).

[Table 2](#) shows how clinicians responded to the survey's list of 8 categories of potential functional additions. The top request was for assistance with selecting the right targeted drug for a given mutation, followed by a request for more links to external knowledge bases, population-level outcome information, and treatment cost information.

The survey results included 35 free-text comments about the app. These comments were coded to 5 categories (navigation, annotation, clinical utility, speed, general) with good interrater agreement ($\kappa = 0.84$). The most common categories were technical aspects of navigation and general comments ([Table 3](#)). Several representative responses to the question of "What features could be improved?" are as follows:

"I think this could be a useful clinical tool if additional features were built in. I would not show this to the average patient that I see, because I generally feel that too much information is confusing to the average pt."

Figure 3: The SMART PCM app allows for user interaction, in order to obtain a quantitative view of the mutation spectrum. In the continued example of a KRAS-mutated lung cancer patient, the user can see that KRAS is the most frequent mutation, and can also see the distribution of other mutations quantified. This information is not displayed in the first visualization (Figure 2) because of the visual clutter.



“Can include more information such as access to my cancer genome. Link to available targeted agents, cost, survival etc.”

“More clear deliniation [sic] of the patient’s findings. All to be legible. Speed may deteriorate when dataset increases. There needs to be a demo to show you how it works. All in all great though!”

Implementing SMART PCM at the ASCO interoperability demonstration

Based upon several of the themes that emerged from the survey results, we implemented additional functionalities. The variant pie visualization was activated so that a user interacting with this data would be brought to the disease- and variant-specific page of My Cancer Genome.³⁰ Additional icons were added to bring the user to disease- or genotype-specific treatment options on HemOnc.org, a collaborative chemotherapy regimen wiki.³¹

The enhanced app was competitively reviewed for inclusion in the American Society of Clinical Oncology (ASCO)’s inaugural interoperability demonstration, and was accepted along with 11 other vendor products.³² SMART PCM received a FHIR-compliant message containing genomic information about a synthetic patient with BRAF-mutated colon cancer in real time from a third-party clinical genomics laboratory, GenoSpace LLC (Cambridge, MA, USA). For the demonstration, the synthetic patient was compared to a locally cached synthetic colon cancer patient population ($N = 415$); visualization construction was performed in real time.

Feasibility of expanding to NGS

Further modifications to the app realized the function of showing data from NGS test results, using synthetic patient data ($N = 150$, see Supplemental Methods). A pie chart was used to stand out gene mutations detected in 1 patient from all gene mutations observed in

similar patients. Further analysis of the distribution of gene mutations was also performed by local outlier factor (LOF) analysis. LOF is an efficient method to identify density-based outliers (see Supplemental Methods).³³ An example is shown in Figure 5 and Figure 6, and Supplementary Figure 2.

DISCUSSION

Cancer is a set of complex diseases whose treatment calls for highly individualized diagnosis and rapidly evolving treatments; indeed, the knowledge base in the genetics of cancer is expanding exponentially.^{34–36} This complexity has been recognized by the broad research and patient care community, and cancer is the first area to be tackled through the Precision Medicine Initiative.¹² As knowledge of the genomic drivers of cancer grows along with the increasing number of druggable targets, matching patients to treatments has become increasingly important.^{2,4} Many oncologists treat a wide variety of cancers, and the potential to move from the existing paradigm of more than 100 distinct types of cancer to thousands or more distinct subtypes is daunting. Given the pace of knowledge generation and the transition to large cancer gene panels and routine whole-exome or whole-genome sequencing, it will be a challenge to incorporate genomic cancer data into existing EHR platforms on a near-real time basis.

We decided that an appropriate initial scope for our clinico-genomic prototype would be to provide information to the clinician for situating a patient relative to other patients in the clinical genomics space. Providing contextually useful patient population comparisons may be a “nice to have” for traditional clinical tests. However, for the complex and highly differentiating results associated with genomic testing, such information is essential for any type of diagnostic/prognostic support, since more active clinical decision rules may depend entirely upon the particular evidence base for the clinical condition of the patient. The informational support provided by the PCM prototype

Figure 4: The SMART PCM app allows for access to external knowledge sources that would otherwise be unavailable to the user through their native EHR system. Shown here is the Gene Wiki page for the gene KRAS, embedded within the app.

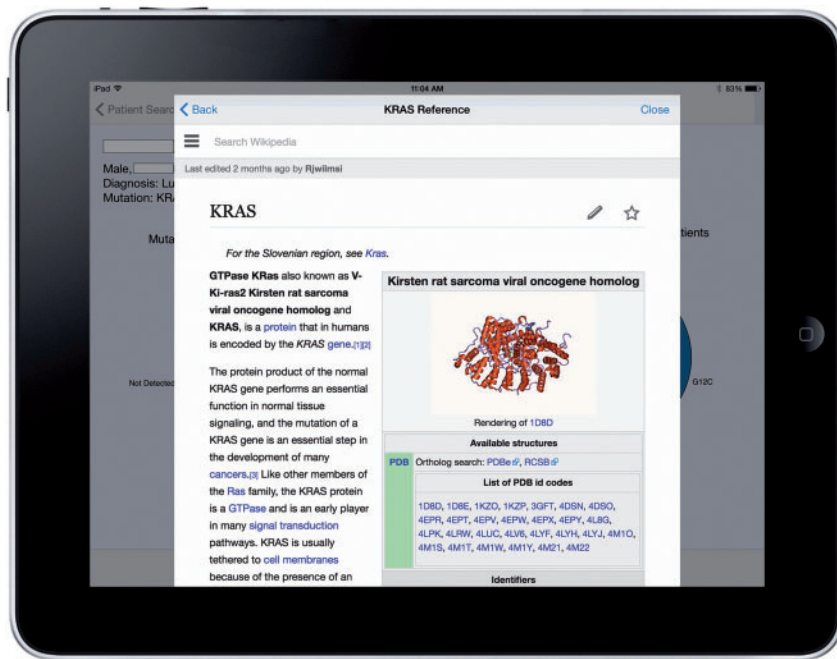


Table 2: Features clinical users would like to see in future PCM designs.

Additional feature	Percent of fellows requesting feature (n = 9) (%)	Percent of attendings requesting feature (n = 5) (%)
Decision support (eg, what drugs will work for my patient?)	67	60
More external knowledge content (eg, links to My Cancer Genome, COSMIC)	44	40
Outcome information (eg, survival)	44	40
Cost information (eg, how much will the drugs cost that could work based on observed mutation?)	33	40
Larger populations (eg, state- or country-level statistics)	0	20
Stratify by criteria such as age, gender, treatment exposure, stage	11	0
Additional visualizations	0	0
All of the above	11	20

Table 3: User comments by category.

Coded user comment category	Comments count
Navigation	8
Annotation	6
Clinical utility	2
Speed	1
General	8

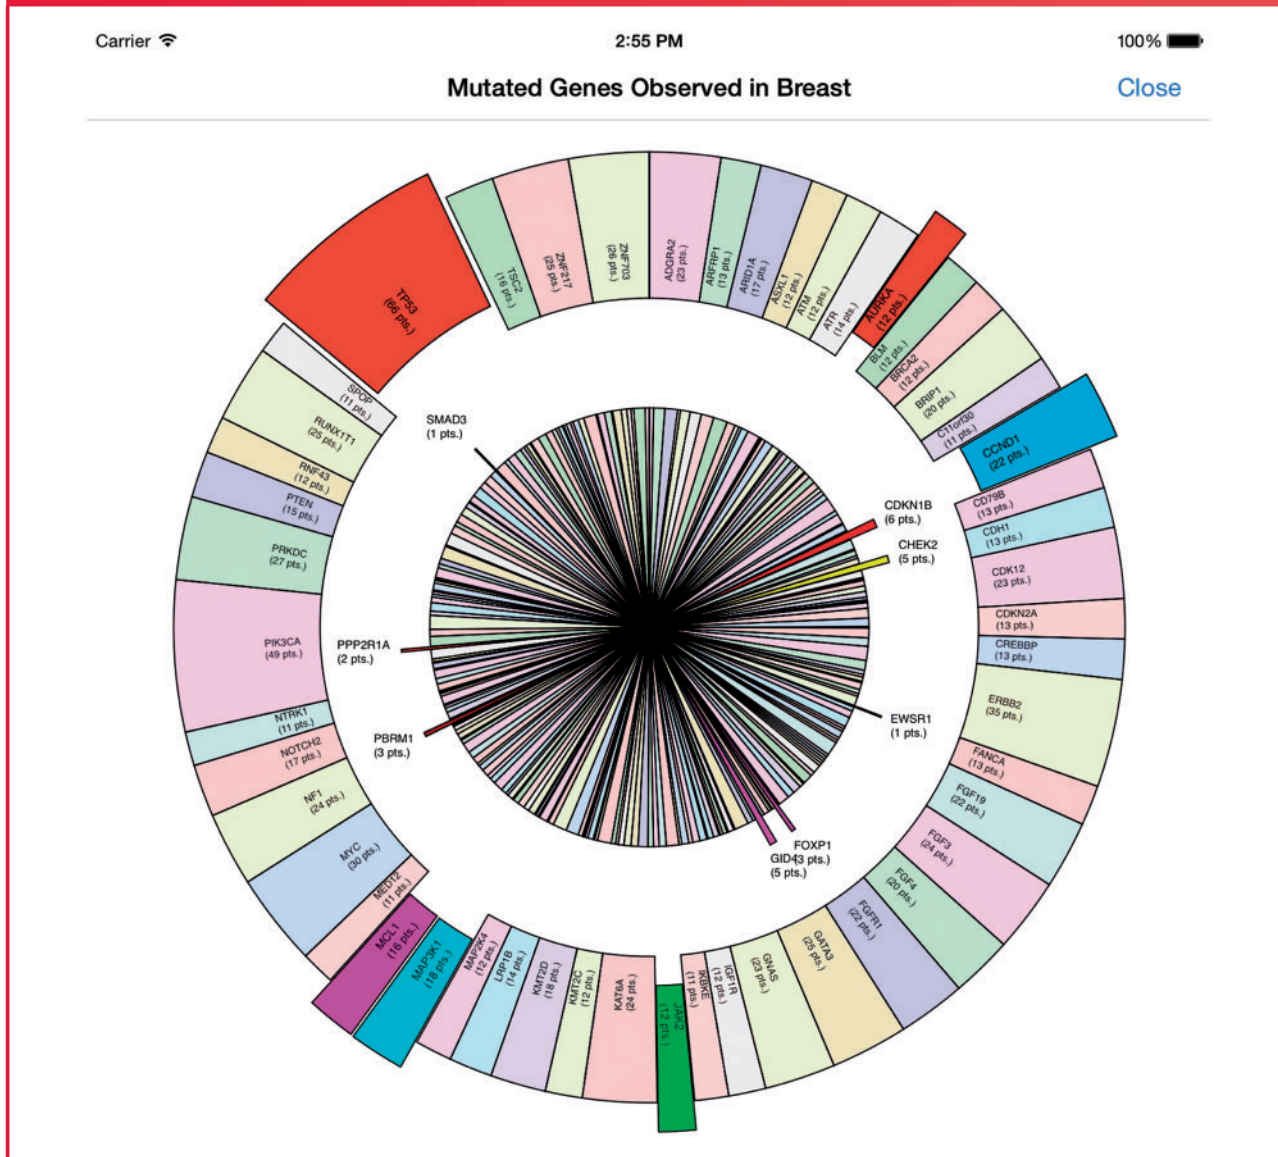
immediately informs the clinician about the patient in a cancer population context and it allows the physician to share—show and tell—this information with the patient.

The PCM app prototype demonstrates how the “app ecosystem” path can keep pace with the underlying medical science.¹³ Our prototype demonstrates how to achieve end-to-end integration with a data

warehouse operating in near–real time with the accompanying EHR system. The SMART on FHIR components enable role-based authentication and authorization for obtaining patient context and population-level data within a firewalled and audited environment. Given the security concerns surrounding PHI, this is a critical aspect. The information required by SMART PCM demonstrates the convenience of the FHIR data model: to satisfy the clinical data requirement, we only had to stand up 4 FHIR Resources: Patient, DiagnosticReport, Specimen, and Observation. Although FHIR had no off-the-shelf way to package the gene and variant data requirements, existing FHIR Resources were readily extended to capture the required information.²¹

By writing the prototype app as a SMART app, we have made it deployable on different EHR systems that have exposed the same FHIR Resources and Extensions. The SMART design also allows accessing patient-specific gene and variant data from a secondary source, such as a commercial DNA sequencing data service. Indeed, a key way in which SMART on FHIR Genomics could stimulate innovation is to offer the right combination of predictable data payloads and secure architectures to simplify “mashing up” of data originating in differing locations, so that capabilities such as those displayed in SMART PCM are no longer exotic.

Figure 5: An example of a synthesized NGS panel result with many detected alterations and 200+ genes with detected alterations in the population.

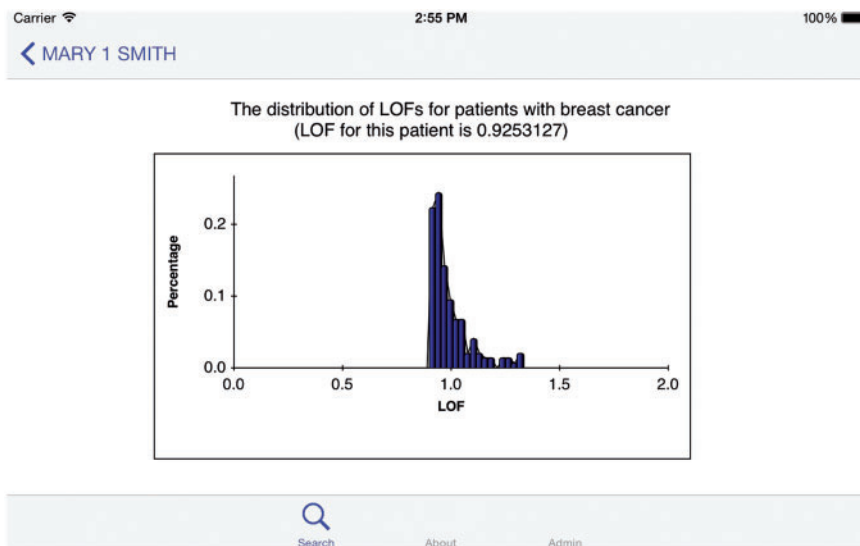


This technical solution does not obviate the need for appropriate data governance policies to ensure that the appropriate level of trust, privacy, and security persists throughout the health care ecosystem. With the growing impetus for post-Meaningful Use interoperability such as The Office of the National Coordinator of Health Information Technology (ONC)'s Federal Health IT Strategic Plan 2015–2020 and embracing of the app concept by some EHR vendors, the time for interoperable standards-based apps is nigh. Through the Argonaut Project (<https://hl7-fhir.github.io/argonauts.html>), an industry-academic consortium, the SMART on FHIR API is being incorporated into 5 major EHR vendor products. The Health Services Platform Consortium, a nonprofit organization with multiple vendor participants, is also promoting the use of apps with a FHIR service layer and EHR integration using SMART. The uptake of SMART by vendors leads to a "win-win" situation where innovation can freely occur and vendors are also free to reuse and adapt open-source app contents into their own products.¹³

Feedback from testers revealed several themes. Foremost is the heterogeneous expectations of clinical oncology users for technology in delivering information. Ranges for the continuous variables in [Supplementary Figure 1](#) averaged 82.75 points out of 100. Additionally, users who rated the app highly with regard to speed or ease of use were not more or less likely to say they would use it clinically. This illustrates a second theme among users: that speed and ease of use must be present, but it is content that will drive the ultimate utility of a clinical app. As demonstrated in [Table 2](#), features such as enhanced decision support and more links to external knowledge bases were the most desired features, whereas additional visualizations and the ability to display data by demographics were less desired. Somewhat surprisingly, only 1 user (out of 14) requested inclusion of larger databases such as COSMIC into the population-level displays. This may reflect the large size of the institutional VUMC database and it is likely

RESEARCH AND APPLICATIONS

Figure 6: The local outlier factor (LOF) distribution for the population; *x*-axis and *y*-axis represent the LOF distribution and the proportion of patients of certain LOF interval in all patients, respectively. The further the LOF value is from 1.0, the more possibility that the genetic mutation observed in that patient is an outlier. The patient represented in Figure 5 (“Mary 1 Smith”) had an LOF of 0.925, suggesting that she is somewhat similar to the other patients in the population.



that implementers with a smaller local database might want external data included to some degree.

As we found in our assessment of coding terminologies (Supplementary Table 1), there are at least 26 structured vocabularies, syntaxes, or Web resources for diseases, 24 for genes, 9 for gene alterations, and 7 for protein alterations. There are also additional vocabularies that have terms explicitly linking 1 or more of these categories (eg, Orphanet links genes to proteins). Although this list may not be comprehensive, choosing from among these was still by no means straightforward. Issues such as version control, provenance, and compatibility must all be considered when selecting terminologies, and for this project we found that the consensus approach was best to reconcile these challenges. However, this process can be time-consuming and does not always guarantee optimal results.

As we demonstrated after the initial user feedback sessions, the addition of more external links is straightforward to implement. External linking capabilities could easily be extended in future work, including patient-centric resources such as Cancer.net and the National Center for Biotechnology Information library of cancer information, and shared decision-making resources such as ClinicalTrials.gov, a clinical trial aggregator. In the absence of an indisputable authoritative knowledge source, multiple links could be provided.

The function of drawing a pie chart showing multiple gene mutations observed in a single specimen substantiates the feasibility of expanding to NGS. However, Figure 5 also demonstrates some significant visualization challenges. First, for the large amount of gene mutations observed, it is hard to integrate all information in one figure. If we display all gene names and their occurrence level, the figure would become too crowded for clinicians to catch things that really matter, while much information would be omitted if we stand out only the most important information. Although alternative visualizations such as bubble plots might overcome some of the weakness of the pie/donut representation, they still suffer from clutter at this scale. Some have advocated circular plots that show connections between

objects or between positions, which is an attractive way to display multidimensional cancer genomics data for scientific publication, but also tends to be over-cluttered for clinical use.³⁷ Even more challenging, there is no consensus on auto-identifying the most valuable information; sometimes high levels of occurrence do not tie in with high importance. For example, the most commonly mutated gene in cancer, tumor protein p53, does not have a therapeutic target. Second, inner connections (eg, physical/genetic interactions) may exist among gene mutations detected; thus, network analysis such as LOF holds the potential to decode potential relationships of mutations including variants of undetermined significance. Visualization of clinico-genomic data is a complex issue that will not be solved for some time. The National Cancer Institute has recognized this complexity and has recently issued a request for applications for Visualization Genomic Data Centers.

CONCLUSION

We built and tested a standards-based clinico-genomic app, with immediate applicability to many cancer patients. The ultimate goals of the SMART PCM app are to (1) provide practitioners with context-dependent population-level cancer mutation information, (2) act as a within-workflow intermediary to select external knowledge bases, and (3) enable a patient-centered and gene-driven shared decision-making model.^{38–40} Patients, caregivers, and clinicians do not wish “precision medicine” to be a mere buzzword, but rather want to know the context of the disease (eg, untreated BRAF-mutated metastatic colon cancer), the prognosis of the disease (eg, the median survival for a patient with untreated BRAF-mutated metastatic colon cancer), and what can be done about the disease (eg, published efficacious treatment regimens for untreated BRAF-mutated metastatic colon cancer). At the same time, personalization does not mean consignment to isolation. If anything, a well-defined niche diagnosis can bring about a sense of solidarity with others so afflicted, as evidenced by the numerous and

strong rare disease coalitions, and sites such as PatientsLikeMe (Cambridge, MA, USA). This can especially be the case when membership in such a group brings the option of targeted therapy, which may also include a unique set of side effects. Apps such as SMART PCM can evolve quickly and nimbly, and offer a new and innovative resource for precision medicine.

Physicians have been described as professional information consumers, and with the increasing power of mobile computing they will increasingly turn to mobile applications for information. These applications will have to be built upon robust data standards to ensure accuracy, speed, and interoperability; however, it will ultimately be content that drives clinical utility and adoption.

ACKNOWLEDGEMENTS

We thank Joseph Burden and Daniel Wenner of the VICC Research Informatics Core for their input into the look and feel of the SMART PCM app; Ariadne K Taylor and Joey Schneider of the VICC Research Informatics Core for their project management; Grahame Grieve, Lloyd McKenzie, and the HL7 FHIR Management Group for their input regarding the modeling of genomic content in FHIR; Amnon Shabo, Mollie Ullman-Cullere, and the HL7 Clinical Genomics work group for their input regarding the representation of genomic content in FHIR; and Nipun Sud and Travis Pittman for their work on the GenoSpace FullView to SMART PCM interface.

COMPETING INTERESTS

The authors have declared that no competing interests exist.

CONTRIBUTORS

J.L.W., D.K., K.D.M., I.S.K., and G.A. conceived the SMART PCM app. D.C., R.O., and L.W. built the SMART PCM app with clinical subject matter expertise provided by J.L.W. J.L.W. and D.C. participated in FHIR Connectathon 7. J.M. hosted the SMART on FHIR server at FHIR Connectathon 7. J.L.W. and M.J.R. conducted the user testing, and M.J.R. carried out the data analysis. S.Z. and H.Y. developed the NGS analysis expansion to SMART PCM. All authors contributed to the initial draft manuscript and approved the final manuscript.

FUNDING

This work was supported by ONC grant 90TR0001/01, the TJ Martell Foundation, the Vanderbilt-Ingram Cancer Center Core Grant (P30-CA68485-18 from NCI), and NLM training grant T15LM7450-12. The funders had no role in study design; in the collection, analysis, and interpretation of data; in the writing of the report; or in the decision to submit the paper for publication.

REFERENCES

1. Van Allen EM, Wagle N, Levy MA. Clinical analysis and interpretation of cancer genome data. *J Clin Oncol*. 2013;31(15):1825–1833.
2. Garraway LA, Lander ES. Lessons from the cancer genome. *Cell*. 2013;153(1):17–37.
3. Masys DR. Biotechnology computing: information science for the era of molecular medicine. *Acad Med*. 1989;64(7):379–381.
4. Levy MA, Lovly CM, Pao W. Translating genomic information into clinical medicine: lung cancer as a paradigm. *Genome Res*. 2012;22(11):2101–2108.
5. Lovly CM, Dahlman KB, Fohn LE, et al. Routine multiplex mutational profiling of melanomas enables enrollment in genotype-driven therapeutic trials. *PLoS One*. 2012;7(4):e35309.
6. Meador CB, Micheel CM, Levy MA, et al. Beyond histology: translating tumor genotypes into clinically effective targeted therapies. *Clin Cancer Res*. 2014;20(9):2264–2275.
7. Su Z, Dias-Santagata D, Duke M, et al. A platform for rapid detection of multiple oncogenic mutations with relevance to targeted therapy in non-small-cell lung cancer. *J Mol Diagn*. 2011;13(1):74–84.
8. Simon R, Roychowdhury S. Implementing personalized cancer genomics in clinical trials. *Nat Rev Drug Discov*. 2013;12(5):358–369.
9. Ashfaq R. Molecular profiling for personalized cancer care. *Clin Exp Metastasis*. 2012;29(7):653–655.
10. Ross JS. Cancer biomarkers, companion diagnostics and personalized oncology. *Biomarkers Med*. 2011;5(3):277–279.
11. Kalia M. Personalized oncology: recent advances and future challenges. *Metabol Clin Exp*. 2013;62 (Suppl 1):S11–S14.
12. Collins FS, Varmus H. A new initiative on precision medicine. *N Engl J Med*. 2015;372(9):793–795.
13. Mandl KD, Mandel JC, Kohane IS. Driving Innovation in Health Systems through an Apps-Based Information Economy. *Cell Syst*. 2015;1(1):8–13.
14. Mandl KD, Mandel JC, Murphy SN, et al. The SMART Platform: early experience enabling substitutable applications for electronic health records. *J Am Med Inform Assoc*. 2012;19(4):597–603.
15. Mandl KD, Kohane IS. No small change for the health information economy. *N Engl J Med*. 2009;360(13):1278–1281.
16. HL7 Version 3: Reference Information Model (RIM). [Accessed August 30, 2015]. http://www.hl7.org/implement/standards/product_brief.cfm?product_id=77
17. Blumenthal D, Tavenner M. The “meaningful use” regulation for electronic health records. *N Engl J Med*. 2010;363(6):501–504.
18. D’Amore JD, Mandel JC, Kreda DA, et al. Are Meaningful Use Stage 2 certified EHRs ready for interoperability? Findings from the SMART C-CDA Collaborative. *J Am Med Inform Assoc*. 2014;21(6):1060–1068.
19. Hardt D. The OAuth 2.0 authorization framework. 2012 [Accessed August 30, 2015]. <https://tools.ietf.org/html/rfc6749>.
20. Sakimura DN, Bradley J, Jones M, de Medeiros B, Jay E. OpenID Connect Standard 1.0-draft 20. 2011. http://openid.net/specs/openid-connect-standard-1_0-20.html.
21. Alterovitz G, Warner J, Zhang P, et al. SMART on FHIR Genomics: facilitating standardized clinico-genomic apps. *J Am Med Inform Assoc*. 2015;22(6):1173–1178.
22. Forbes SA, Tang G, Bindal N, et al. COSMIC (the Catalogue of Somatic Mutations in Cancer): a resource to investigate acquired mutations in human cancer. *Nucleic Acids Res*. 2010;38(Database issue):D652–D657.
23. Gray KA, Yates B, Seal RL, Wright MW, Bruford EA. Genenames.org: the HGNC resources in 2015. *Nucleic Acids Res*. 2015;43(Database issue):D1079–D1085.
24. Sioutos N, de Coronado S, Haber MW, Hartel FW, Shaiu WL, Wright LW. NCI Thesaurus: a semantic model integrating cancer-related clinical and molecular information. *J Biomed Inform*. 2007;40(1):30–43.
25. Farrell CM, O’Leary NA, Harte RA, et al. Current status and new features of the Consensus Coding Sequence database. *Nucleic Acids Res*. 2014;42(Database issue):D865–D872.
26. Good BM, Clarke EL, de Alfaro L, Su AI. The Gene Wiki in 2011: community intelligence applied to human gene annotation. *Nucleic Acids Res*. 2012;40(Database issue):D1255–D1261.
27. Faulkner L. Beyond the five-user assumption: benefits of increased sample sizes in usability testing. *Behav Res Methods, Instruments, Comput*. 2003;35(3):379–383.
28. Danciu I, Cowan JD, Basford M, et al. Secondary use of clinical data: the Vanderbilt approach. *J Biomed Inform*. 2014;52:28–35.
29. Wattanasin N, Porter A, Ubaha S, et al. Apps to display patient data, making SMART available in the i2b2 platform. *AMIA Annu Symp Proc*. 2012;2012:960–969.
30. Swanton C. My Cancer Genome: a unified genomics and clinical trial portal. *Lancet Oncol*. 2012;13(7):668–669.
31. Warner JL, Cowan AJ, Hall AC, Yang PC. HemOnc.org: A Collaborative Online Knowledge Platform for Oncology Professionals. *J Oncol Pract*. 2015;11(3):e336–e350.
32. Krauss JC, Warner JL, Maddux SE, et al. Data sharing to support the cancer journey in the digital era. *J Oncol Practice*. 2015 (In press).

33. Breunig MM, Kriegel H-P, Ng RT, Sander J. LOF: identifying density-based local outliers. *ACM sigmod record*; 2000: *ACM*; 2000: 93-104.
34. Schilsky RL. Implementing personalized cancer care. *Nat Rev Clin Oncol*. 2014;11(7):432–438.
35. Rieth MJ, Osterman TJ, Warner JL. Advances in website information resources to aid in clinical practice. *American Society of Clinical Oncology educational book / ASCO American Society of Clinical Oncology Meeting*. 2015;35:e608–e615.
36. Warner JL. Grappling with the data explosion in oncology. *Oncol Hematol Rev*. 2015;11(2):102–103.
37. Schroeder MP, Gonzalez-Perez A, Lopez-Bigas N. Visualizing multidimensional cancer genomics data. *Genome Med*. 2013;5(1):9.
38. Jekunen A. Clinicians' expectations for gene-driven cancer therapy. *Clin Med Insights Oncol*. 2014;8:159–164.
39. Abernethy A, Abrahams E, Barker A, et al. Turning the tide against cancer through sustained medical innovation: the pathway to progress. *Clin Cancer Res*. 2014;20(5):1081–1086.
40. Bombard Y, Bach PB, Offit K. Translating genomics in cancer care. *J Natl Compr Canc Netw*. 2013;11(11):1343–1353.

AUTHOR AFFILIATIONS

¹Department of Medicine, Division of Hematology and Oncology, Vanderbilt University, Nashville, TN, USA

²Department of Biomedical Informatics, Vanderbilt University, Nashville, TN, USA

³Boston Children's Hospital Computational Health Informatics Program, Boston, MA, USA

⁴Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA

⁵Independent Consultant, New York, NY, USA

⁶Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA

⁷Department of Pediatric Endocrinology, Boston Children's Hospital, Boston, MA, USA

⁸Vanderbilt-Ingram Cancer Center, Vanderbilt University Medical Center, Nashville, TN, USA

⁹Department of Electrical Engineering and Information Science, University of Science and Technology of China, Hefei, China

¹⁰Department of Life Science, University of Science and Technology of China, Hefei, China