Modeling the costs of clinical decision support for genomic precision medicine

Patrick C. Mathias, MD, PhD[1], Peter Tarczy-Hornoch, MD[2], Brian H. Shirts, MD, PhD[1]
[1]Department of Laboratory Medicine, [2]Departments of Biomedical Informatics and Medical Education, Pediatrics, and Computer Science and Engineering; University of Washington, Seattle, WA

Abstract

Clinical decision support (CDS) within the electronic health record represents a promising mechanism to provide important genomic findings within clinical workflows. To better understand the current and possible future costs of genomic CDS, we leveraged our local CDS experience to assemble a simple model with inputs such as initial cost and numbers of patients, rules, and institutions. Our model assumed efficiencies of scale and allowed us to perform a one-way sensitivity analysis of the impact of each model input. The number of patients with genomic results per institution was the only single variable that could decrease the cost of CDS per useful alert below projected genomic sequencing costs. Because of the prohibitive upfront cost of sequencing large numbers of individuals, increasing the number of institutions using genomic CDS and improving the efficiency of sharing CDS infrastructure represent the most promising paths to making genomic CDS cost-effective.

Introduction

Genomic information has potential to dramatically increase the complexity of test result delivery. For optimal delivery of precision medicine, physicians will need to be aware of genetic results from testing performed years earlier. Specific genomic variants may alter clinical care for only a small fraction of patients, and long-term awareness of genomic sequencing results could put an enormous burden on physicians unless the current mode of clinical genetic information delivery is altered. Current an electronic health record (EHR) systems are designed for communication of on-demand laboratory results, yet many important results are not acted upon by clinicians for a variety of reasons, including failure to retrieve results and incorrect interpretation of results[1]. Furthermore, a majority of institutions in one recent survey indicated that at least a subset of genomic information in their EHRs is only available in unstructured formats[2]. Therefore, active clinical decision support (CDS) alerts coupled with structured genomic result information within EHRs have been proposed as an ideal mechanism to alert clinicians of pertinent genomic findings by incorporating genomic data into clinical workflows[3].

As a component of the New Exome Technology in Medicine study, CDS rules for genomic results were designed and implemented in one of our institution's EHRs[4]. After completing this work, we noted there was limited existing literature exploring the costs of effectively communicating genomic information over long periods of time. Recent work has demonstrated that reporting incidental findings can be cost effective in several situations, but assumed that the costs of sequencing were $500 and did not factor in expenses associated with just-in-time result delivery and CDS[5]. Because the costs of EHR configuration for genomic result display and CDS have not been previously described, we modeled potential costs of design, implementation, and maintenance of CDS for genomic medicine under a variety of scenarios.

Methods

To estimate the costs of employing genomic CDS, we performed a one-way sensitivity analysis using cost per useful CDS alert as the primary evaluation metric and made a number of assumptions based on literature and experience designing and implementing genomic CDS at our institution. We modeled the effect of 8 different parameters, each of which was varied while fixing the other 7 to our estimates for the current state for our institution, assuming that the cost of developing each subsequent CDS rule after the first is a fixed proportion of the previous rule, until a minimum cost is reached, to allow for efficiency of scale. The same efficiency assumption is applied to distributing rules between institutions, with the recognition that the current state of distinct EHR builds makes it difficult to share rules. The following equation was used to model the impact of each variable (abbreviations listed in Table 1):

$$Cost/alert = \frac{C * (1 + M * t)}{n_{patients} * n_{rules} * P_{benefit}} * \sum_{r=1}^{n_{rules}} \max\left(Eff_{rule}^{r-1}, floor_{rule}\right) \sum_{i=1}^{n_{inst}} \max\left(Eff_{inst}^{i-1}, floor_{inst}\right)$$

where t indicates the average amount of time before a useful alert is triggered, $floor_{rule}$ refers to the minimum proportion of initial rule cost that a subsequent rule may cost, and $floor_{inst}$ refers to the minimum proportion of initial rule cost that a separate institution would have to spend to develop a rule.

At our institution we built 8 CDS rules into our EHR as part of a clinical sequencing exploratory research grant and used this experience to develop ranges of inputs for the current state of genomic CDS. By accounting for time and labor spent developing our CDS rules locally, we estimated that the costs of designing, building, and implementing a single genetic CDS rule at a single institution would be between $4000 and $7500 of effort divided between EHR decision support architects, laboratory information specialists, and institutional genetics committees. CDS rules must be maintained over this time, and since the IT industry often accounts for a maintenance cost of 20% of the initial cost per year, we assumed maintenance would be 10-25% of implementation costs per year for 30 years, which would be the average time before a rule is used. Because events that require genomic information can occur anytime during a patient's lifetime, the lifetime costs of maintaining CDS rules for a population of patients, or regularly rebuilding rules as EHR and laboratory systems change, will reasonably be 3 to 7.5 times the cost of initial rule implementation.

The CDS rules were put in place at our institution for an initial cohort of 500 patients, and we modeled a range of 200 patients to 41,250 patients per institution, deriving the high estimate from 75% of the average patient population of 55,000 per institution being screened with genomic testing. We have implemented 8 distinct CDS rules at our institution and reasoned that at a maximum there may be a few hundred genomic variants that would warrant genomic CDS rules. A survey of Clinical Sequencing Exploratory Research sites revealed 2 sites with structured genomic CDS[6], but we assumed that additional sites have likely implemented this since that time and thus used 10 sites as the current state for number of institutions. From known allele and disease frequency, we estimated that the average rule might benefit approximately 1% of patients, reasoning that although any actionable pharmacogenomic variant may be observed in >90% of patients[7], only a small proportion of patients with each specific variant would be prescribed a relevant drug in their lifetime.

Our experience implementing rules indicated that adding additional CDS rules required less work than the first rule, and we estimated that each rule took about 80% of the effort of the previous rule but modeled a relatively narrow range of rule efficiencies (70-85%) since some of the effort in implementing rules, such as testing, is unlikely to be avoided. We modeled a slightly larger efficiency range for collaborating institutions (70-100%), but assumed that the current state efficiency between institutions is 99% since rules cannot currently be easily translated between different EHR builds. We made relatively optimistic assumptions about the efficiency floor for increasing the number of rules: the minimum cost for a rule is 10% of the cost of the initial rule. We made a similarly optimistic assumption for sharing rules between institutions: the minimum cost for a new institution for a shared rule is at least 5% of the cost of the initial CDS rule.

Table 1. Variables modeled in one-way sensitivity analysis, with ranges for inputs and resulting projected costs per useful alert.

	Input Variables		
Variable	**Low**	**Current**	**High**
Initial cost for single CDS rule development and implementation C	4000	6600	7500
Number of CDS rules n_{rules}	3	8	300
Number of institutions n_{inst}	2	10	500
Average probability each rule will benefit a patient $P_{benefit}$	0.0025	0.01	0.02
Patients with sequence per institution $n_{patients}$	200	500	41250
Rule maintenance rate per year (as a proportion of development costs) M	0.1	0.2	0.25
Efficiency gained for multiple CDS rules Eff_{rule}	0.7	0.8	0.85
Efficiency gained by collaborating institutions Eff_{inst}	0.7	0.99	1

By integrating current state assumptions with the total cost for developing the 8 CDS rules implemented at our institution, our best estimate of our initial rule creation and testing cost was approximately $6600, which we used in our model as the current state estimate of initial CDS rule cost.

Results

After applying our current state inputs for each of the variables, we calculate the lifetime cost of communicating genomic information to a physician at the point of care to be $4600 per genomic decision support event for 8 scenarios currently covered by CDS rules. This estimate assumes that rules are maintained for the lifetime of the patients tested and the duration before an alert is triggered is 30 years.

Figure 1 presents the results for our one-way sensitivity analysis, illustrating the variables that most influenced model outcomes. The average probability a rule would benefit each patient (range: 0.25-2%), the number of patients with genomic sequence stored at an institution (range: 200-41250), the number of institutions sharing CDS rules (range: 2-500), the number of CDS rules built (range: 3-300), and the potential efficiency gained by collaborations to build CDS that are usable across institutions (range: 70-100% cost of previous rule for each successive rule) all showed potential to significantly decrease the cost of CDS. However, only having a large number of patients with genomic sequences at each institution decreased costs per alert below $1000.

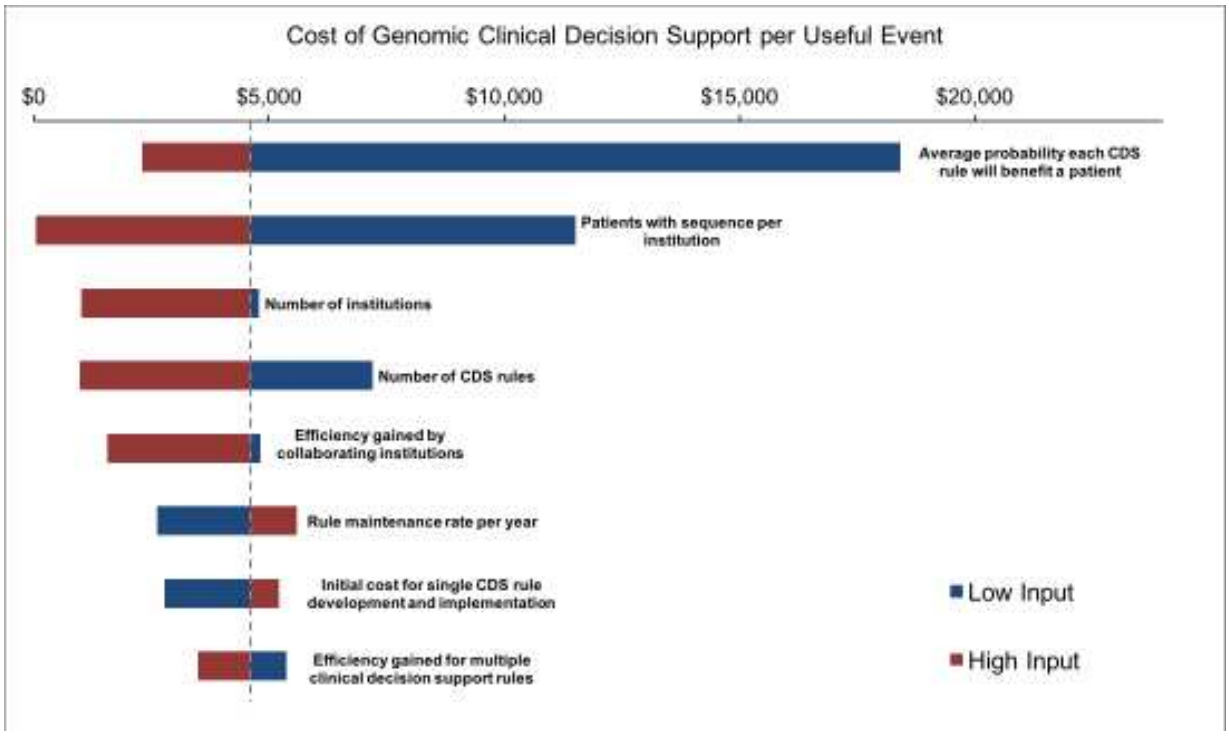

Figure 1. Results of one-way sensitivity analysis characterizing the effects of different ranges of inputs on the cost of CDS per useful event.

Discussion

To better understand the costs of implementing genomic CDS, we developed a simple model to gauge the relative importance of 8 variables on the cost per useful CDS alert. Because of the limited literature on genomic CDS costs, we made assumptions for reasonable inputs based on our local experience developing CDS. Our model relies on the leveraging efficiencies of scale in implementing CDS, both within an institution and between institutions. Based on our assessment of the current state at our institution, we estimate that the cost per useful genomic CDS alert is $4600, which is almost an order of magnitude greater than the cost of genomic sequencing used in current cost-benefit analyses. Some of these variables we evaluated are closely related. The probability that a rule will benefit a patient is inversely correlated with the number of CDS rules; after clearly useful rules are built for common variants and common situations, then additional rules will necessarily be for less common events. CDS rules for rare

variants benefit fewer individuals, but have similar development and maintenance costs. Several different strategies to efficiently implement and maintain multiple CDS rules could decrease CDS costs per patient receiving benefit. Our analysis shows that efficiencies of scale gained from sharing CDS rule implementation across many institutions or implementation of genomic testing in a greater proportion of patients would substantially reduce per patient genomic CDS costs.

The only single factor that could reduce the cost of CDS per useful event to below several hundred dollars was increasing the number of patients with genomic results to over 75% of the population, which would take an extraordinary initial cost and effort. Another option to reduce the per-event cost of CDS to a manageable level is to address multiple variables simultaneously, such as increasing cooperation between institutions for building CDS rules, increasing the number of meaningful CDS rules, and improving the ease and efficiency of the process to build genomic decision support in current EHR systems. For example, increasing efficiency to reduce the cost of sharing between institutions to 70% of initial costs for every additional institution sharing CDS rules in addition to sharing between 500 institutions together could drive cost costs per alert to less than $300. Such an approach requires broad agreement and distribution of useful CDS rules, and emphasizes the importance of the Agency for Healthcare Research and Quality Patient Centered Outcomes Research CDS initiative[8] and the efforts of the Institute of Medicine Displaying and Integrating Genetic Information Through the EHR (DIGITizE) working group[9]. Additionally, accepted CDS rules must be combined with data exchange standards and true interoperability of genomic data in order to lower costs and increase efficiency, which the Office of the National Coordinator for Health Information Technology has strongly encouraged but will likely require considerable effort from the government, vendors, and institutions to make a reality. The Fast Health Interoperability Resources (FHIR) standard offers a promising solution to tackle the problem of interoperability, and applications using genomic data have already been demonstrated using this standard[10]. However, considerable further effort will be required to apply this technology across diverse settings.

The primary limitation of our study is our reliance on model inputs derived from our local experience implementing genomic CDS. It is unclear whether our experience is representative of other institutions. However, given the lack of interoperability between EHRs, the general processes of genomic result formatting and rule creation and testing are unlikely to be dramatically different between institutions and EHRs. Another limitation of our study is that our model relies on efficiencies of scale within and between institutions. To our knowledge, there are no detailed studies exploring the scalability of genomic CDS to date, so we built our model with relatively optimistic assumptions about scalability. For example, we set the minimum cost of any additional rule added to 5% of the cost of the original rule. In reality, the effort of testing alone would likely exceed this cost, yet the model demonstrates that, even with optimistic inputs, decreasing the cost per alert may be challenging.

Although the costs of performing genomic sequencing have decreased dramatically, our modeling indicates that the cost of effectively communicating genomic information to clinicians with current reporting systems is substantial and should be factored into any discussion of cost-effectiveness. President Obama's Precision Medicine Initiative proposes extensive work to improve our understanding of the genomic underpinnings of disease, yet only 2% of this effort is proposed for information technology work, which will be critical to integrate precision medicine with existing healthcare delivery systems. Additional inter-institutional efforts to build efficient, scalable systems for communicating genetic test results and actionable interpretations at the point of care are needed to decrease implementation costs of decision support to effectively communicate pertinent genomic information at the point of care.

Conclusion

Using the local experience of building genomic CDS at our institution, we demonstrate a straightforward model for assessing the costs of CDS per useful alert. Our data suggests that building genomic CDS locally with current EHR infrastructure incurs per alert costs that are higher than anticipated costs of genomic sequencing. While sequencing large numbers of patients could decrease the cost to less than $100 per alert, the upfront cost of doing so would be tremendous. Either increasing the number of institutions with genomic CDS or improving the efficiency of sharing CDS in isolation could decrease costs per alert, but combining these approaches likely offers the most efficient path to building alerts cost-effectively.

Acknowledgements: We would like to thank Aiden Garver-Hume, Joe Smith, and Chuck Rohrer for their efforts in helping to implement genomic CDS and feedback on the effort required to do so. We would also like to thank all our colleagues in the Precision Medicine Informatics Group, especially Sean Mooney and Beth Devine, for their valuable feedback.

Financial Support: This work was funded by the National Human Genome Research Institute, the National Center for Advancing Translational Sciences, and the National Cancer Institute through U01 HG006507, UL1TR000423, and U01 HG007307.

References

1. Gandhi TK, Kachalia A, Thomas EJ, et al. Missed and delayed diagnoses in the ambulatory setting: a study of closed malpractice claims. Ann Intern Med. 2006;145:488-496.
2. Shirts BH, Salama JS, Aronson SJ, et al. CSER and eMERGE: current and potential state of the display of genetic information in the electronic health record. J Am Med Inform Assoc [Internet]. 2015 Jul 3 [cited 2015 Sep 24]. Available from: http://jamia.oxfordjournals.org/content/early/2015/07/02/jamia.ocv065.long
3. Welch BM, Eilbeck K, Del Fiol G, Meyer LJ, Kawamoto K. Technical desiderata for the integration of genomic data with clinical decision support. J Biomed Inform. 2014;51:3-7.
4. Nishimura AA, Shirts BH, Dorschner MO, et al. Development of clinical decision support alerts for pharmacogenomic incidental findings from exome sequencing. Genet Med [Internet]. 2015 Mar 5 [cited 2015 Sep 24]. Available from: http://www.nature.com/gim/journal/vaop/ncurrent/full/gim20155a.html
5. Bennette CS, Gallego CJ, Burke W, Jarvik GP, Veenstra DL. The cost-effectiveness of returning incidental findings from next-generation genomic sequencing. Genet Med. 2015;17:587-595.
6. Tarczy-Hornoch P, Amendola L, Aronson SJ, et al. A survey of informatics approaches to whole-exome and whole-genome clinical reporting in the electronic health record. Genet Med. 2013;15:824-832.
7. Van Driest SL, Shi Y, Bowton EA, et al. Clinically actionable genotypes among 10,000 patients with preemptive pharmacogenomic testing. Clin Pharmacol Ther. 2014;95:423-431.
8. Initiative to Disseminate, Implement PCOR Findings via CDS [Internet]. Rockville, MD: Agency for Healthcare Research and Quality; 2015 July [cited 2015 Sep 22]. Available from: http://www.ahrq.gov/news/pcor.html
9. DIGITizE: Displaying and Integrating Genetic Information Through the EHR [Internet]. Washington, DC: Institute of Medicine; 2015 [cited 2015 Sept 22]. Available from: http://iom.nationalacademies.org/Activities/Research/GenomicBasedResearch/Innovation-Collaboratives/EHR.aspx
10. Alterovitz G, Warner J, Zhang P, et al. SMART on FHIR Genomics: Facilitating standardized clinic-genomic apps. J Am Med Inform Assoc [Internet]. 2015 Jul 21 [cited 2015 Sep 22]. Available from: http://jamia.oxfordjournals.org/content/early/2015/07/21/jamia.ocv045