

Structural Patterns in Chronic Disease Clinical Practice Guidelines Formalized for Clinical Decision Support

Samson W. Tu MS^{1,2}, Tanya Podchiyska MS¹, Connie Oshiro PhD¹, Susana Martins MD MSc^{1,3}, Michael Ashcraft MD¹, Justin Chambers¹, Amy Robinson PharmD^{1,3}, Paul Heidenreich MD MS^{1,2}, Mary K. Goldstein MD MS^{1,2}

¹VA Palo Alto Health Care System, Palo Alto, CA; ²Stanford University, Stanford, CA;

³VA Program Evaluation and Resource Center, Office of Mental Health and Suicide Prevention, Menlo Park, CA

Abstract

Making recommendations from clinical practice guidelines (CPGs) computable for clinical decision support (CDS) has typically been a laborious and costly process. Identifying domain-specific regularities helps clinicians and knowledge engineers conceptualize, extract, and encode evidence-based recommendations. Based on our work to provide complex CDS in the management of multiple chronic diseases, we propose nine chronic disease CPG structural patterns, discuss considerations in representing the necessary knowledge, and illustrate them with the solutions that our CDS system provides.

Introduction

Professional and healthcare organizations encourage the use of clinical practice guidelines (CPGs) to promote evidenced-based best practices, improve patient well-being, and reduce cost. Computer-interpretable guidelines (CIG) formalize narrative CPG recommendations to expedite the adoption of CPG recommendations.¹

Formalizing CPGs into CIGs continues to be slow and costly. Past work on the encoding process focused on implementation steps to bridge the gap between narrative CPGs and formalized CIGs by systematizing the steps,² by creating levels of representations,³ by creating standardized and shareable representations,⁴ and by enabling the specification of computable guidelines to occur concurrently with guideline development.⁵

Some researchers have investigated the orthogonal problem of conceptualizing structural patterns that characterize guideline recommendations in different clinical domains. Peleg and Tu⁶ defined computable templates for screening and immunization guideline recommendations. Kaiser and Marcos⁷ analyzed CPGs to represent their procedural knowledge in terms of workflow patterns. Moser and Miksch⁸ created simple structural and temporal patterns from textual analysis of guidelines. Similarly, Gooch and Roudsari⁹ presented a literature review to identify challenges involved in developing process-oriented health information systems and to summarize approaches to developing such systems. Abstracting structural patterns of CPGs may facilitate not only knowledge synthesis and encoding for CDS, but also guideline standardization and integration. Common structural patterns of CPGs may ultimately facilitate the development of CDS tools that take advantage of them.

Our goal is to systematize structural patterns common to chronic disease CPGs by leveraging our experiences in developing CDS systems (CDSSs) based on such guidelines and to present a categorization of the challenges involved in their conceptualization and encoding. These structural patterns present not only technical issues in knowledge representation, but also a cognitive conceptualization of distinctive features in chronic disease management. They are not exactly “design patterns” defined in software engineering,¹⁰ as we do not present solution templates. Instead, we will describe, as exemplars, the approaches we have taken in representing these structural patterns. We will discuss considerations for designing a solution and the consequences of implementing such a solution. To our knowledge, the only similar prior work on patterns of chronic disease guidelines is the “scenario-based” approach of the PRODIGY III system in the United Kingdom, where the entire guideline representation was designed to facilitate the encoding of chronic disease guidelines.¹¹ Our guideline representation borrows from and extends that system.

Background

As a guide to best practices for healthcare providers, CPGs are intended to optimize patient healthcare and health through facilitating the dissemination of up-to-date, evidence-based care recommendations and reducing unwarranted variability in healthcare practice. Chronic disease guidelines are especially relevant today as the prevalence of chronic diseases is rising and many patients have multiple chronic conditions.¹² Chronic care addresses long-term illness and is often complex, requiring serial assessments of patient health status and risk factors intended to inform appropriate and timely interventions to achieve care goals and meet patient preferences. Strategies for goal setting and treatment,

such as stepped care and medication management—from addition, substitution, and stoppage of medications to dose adjustment—are important tools available to clinicians.

For more than ten years, our group at Veterans Affairs (VA) Palo Alto Health Care System has used ATHENA-CDSS, a knowledge-based system to provide CDS for guideline-based care, to investigate issues related to the operationalization, testing, and deployment of chronic disease CPGs.¹³⁻¹⁵ Initially focused on hypertension (HTN), ATHENA-CDSS has been extended to the operationalization of guidelines for hyperlipidemia, chronic kidney disease (CKD), type-2 diabetes mellitus (T2DM), and heart failure (HF).^{15,16} ATHENA-CDS knowledge engineers use the Protégé¹⁷ tool to formalize guideline recommendations into CIG in terms of the EON guideline model.¹⁸⁻²⁰ The EON execution engine applies the resulting knowledge bases to electronic health record (EHR) data to generate patient-specific guideline-based recommendations for the patients.²¹

In our Medication Safety (MedSafe) CDSS—the successor of ATHENA-CDSS—we leveraged VA’s Veterans Integrated Services Network (VISN) 21 Performance Measure (PM) Dashboard and integrated population-based clinical assessment and individualized guideline-based CDS for patients with hypertension, T2DM, and CKD. On the Dashboard, clinician users view the patients not meeting clinical goals as a percentage of their overall patient population. They can drill down to view a patient’s performance measure status, relevant clinical data, and actionable, patient-specific MedSafe-CDS recommendations based on relevant CPGs.

The EON guideline model structures a CIG in terms of eligibility criteria; goals (e.g., target blood pressures for different subpopulations); drug usage properties (e.g., first-line drug, strong contraindications, maximum dosage),²⁰ a clinical algorithm consisting of sequenced and parallel nodes of patient scenarios, decisions (also known as “choice steps”), and actions. A medical concept hierarchy provides the vocabulary for encoding guideline recommendations. An encoded guideline may include scenario-specific subguidelines or it may include subguidelines that are executed sequentially independent of the containing guideline. In this paper, we formulate structural patterns that we have identified in chronic disease guidelines and document our approach to modeling them in MedSafe-CDSS and ATHENA-CDSS using constructs of the EON guideline model.

Method

Over the years our project maintained a manual for knowledge engineering and “Rules documents” that abstract, synthesize, and structure relevant recommendations from guidelines and from available literature, as detailed in the implementation model described in Tso et al.²² One of the co-authors (SWT) documented draft structural patterns as they emerged through meetings, email discussions, and implementations of MedSafe-CDS guideline knowledge bases. The other co-authors reviewed them and suggested consolidations and revisions.

Results

Table 1 summarizes the structural patterns, organized into three large categories, with short definitions and examples. We describe these patterns in detail in the rest of the section.

Table 1. Summary of Structural Patterns of Chronic Disease Guidelines.

Appropriateness of guideline, and ultimately CDS, recommendations		
Scoping applicable patient cases	Determine which patients are in scope based on applicability of implemented recommendations	e.g., Ruling out complex hypertension cases involving four or more medications
Data availability	Missing critical data needed for making specific recommendations	e.g., Lack of recent blood pressure data
Therapeutic strategies		
Applying therapeutic goals	Having multiple, possibly interacting, goals to achieve in the management of one or more conditions	e.g., Glycemic control, preservation of kidney function, and prevention of eye and foot complications for a diabetic patient
Recognizing patient scenarios	How to define high-level patient states that help to structure management decisions	e.g., Whether the patient has been prescribed no, one, two, or three or more antihypertensive agents

Assessing current treatment	How to model and display messages about the current treatment	e.g., Recheck the GFR value when it is between 45 and 60 ml/min/1.73 m ²
Determining degree of detail in therapeutic recommendations	Recommending only the most optimal therapeutic choice(s) versus showing all applicable choices in a stratified list	e.g., Recommending only metformin or including pros and cons of second-line drug such as empagliflozin
Drug information representation		
Drug usage information	How to model and use drug properties that affect their usage	e.g., indications, contraindications, possible side effects
Preference over drug classes and specific drugs	How to model and use guideline-based or institutional drug preferences	e.g., drugs that are on formularies
Constraint on drug selections	How to model and use information on when drugs should or should not be used together	e.g., Only one ACE inhibitor should be used to manage a patient's hypertension

1. Appropriateness of guideline and CDS recommendations.

The appropriateness of guideline and CDS recommendations depends on two main factors: whether the encoded recommendations in the CIG are applicable to a patient, and whether there are sufficient data to generate meaningful recommendations.

Scoping applicable patient cases

Every CDSS that implements the recommendations of a guideline must identify patient cases for which the guideline is applicable. There may be situations that are too unusual or where effective management is too complex or where data requirements are too extensive. Often it is possible to formalize the definition of the population that is targeted by the guideline and implemented in CDSS as a set of formal eligibility criteria. In some cases, the CDSS will need to explain why a case is out of scope and point to other useful resources. The eligibility criteria can be evaluated by the CDSS or by some external filtering service, with particular attention to whether lack of information (e.g., missing laboratory test results) rules in or rules out a patient case.

In MedSafe-CDS encoding, we have choice to either (1) stipulate a set of Boolean criteria all of which should evaluate to 'true' or 'unknown' for a patient to be eligible for CDS or (2) create a scenario in the clinical algorithm whose precondition is a disjunction of all situations that are out of scope for the CDS. In the latter case, subsequent action nodes can then deliver specific messages for each situation and refer to appropriate resources. For example, if the out-of-scope conditions for the CDSS's implementation of T2DM guideline are (1) use of insulin, (2) use of non-encoded drugs, and (3) use of three or more oral hypoglycemic agents, then a scenario is created whose precondition is a disjunction of the three criteria. If the three criteria may be true independently of each other, then one choice step may be used to set up a non-exclusive OR branching point. If they are mutually exclusive, then the rule-in criteria of the three actions need to be ascertained as appropriate. If there is an ordering among them, a branching ladder can be used to sequence them.

The consequences for the choices are that, if implemented as eligibility criteria that filter out non-applicable cases (choice 1 above), the execution engine has fewer cases to process, but the system can neither explain the filtering process nor generate messages about possible resources that may be relevant, which choice 2 allows. However, users not being shown any CDS information about excluded patients may reduce alert fatigue. Thus, the eligibility criteria are a relatively easy way to filter out patient cases and can reduce the cost of displaying explanations and pointers to other resources. A clinical advisory group for the project, consisting of front-line clinician representatives of the target clinician population of CDS users, can provide input to help determine the balance between providing resources for unusual cases versus the extra alerts that are not directly making a recommendation.

Lack-of-Critical-Data Pattern

Disease management often requires up-to-date information, such as blood pressure readings for the management of hypertension, to determine which recommendations apply. To assist users, the CDSS message should clearly display

which data element is missing and possibly show the most recent value of the data element, especially if it is not in the required range.

A commonly occurring scenario in clinical care is the patient whose laboratory test results are several months old. It is usually a matter of clinical judgment rather than CPG statement to determine if the laboratory values are current enough for clinical decision making. In some cases, a laboratory test is done on the day of the visit and its result may not be available at the time the CDS recommendations are computed. Blood pressure measurements may not yet be entered into electronic records at the time recommendations are computed. In the MedSafe-CDS knowledge base, we sometimes choose to create a scenario defined by the lack of critical data and create decision and action nodes to determine which and possibly when data collection actions are needed. The primary CDS outputs consist of messages explaining what critical data are missing and recommendations to order appropriate tests and measurements. In an earlier implementation of the hypertension guideline,²³ if the CDSS did not have the requisite blood pressure data for a patient, it made alternative assumptions on the blood pressure status, and generated two separate sets of CDS recommendations on the management of hypertension for the patient. At a more granular level, whether a drug should be recommended may depend on certain laboratory test values being in appropriate ranges. In that case, the MedSafe-CDS makes the drug recommendation conditional on the missing data.

By explicitly taking the availability of critical data into account, a CDSS can inform the clinicians of the data gap and possibly the need to order necessary tests.

2. Therapeutic strategies

In the category of therapeutic strategies, we distinguish four patterns: (1) applying treatment goals, (2) recognizing the scenario that is appropriate for a patient, (3) assessing the patient's current treatment, and (4) determining appropriate drug treatment.

Applying therapeutic Goals

For the signature diseases of the guidelines, often there are recommendations to achieve and maintain certain therapeutic targets. The specific target to achieve and maintain may be different for subpopulations of patients. The EON guideline model allows us to encode, for each guideline, a collection of if-then statements that set appropriate targets for subpopulations in the scope of the CDSS.

Additionally, for a particular disease, there may be multiple goals to achieve. For example, an HTN guideline may include not only a target blood pressure but also the goal of secondary prevention of cardiovascular disease. A diabetes guideline may be concerned about glycemic control, maintenance of kidney function, and prevention of eye and foot complications. For a patient with multimorbidity, for example one with diabetes, hypertension, and chronic kidney disease, a clinician necessarily needs to take multiple treatment goals into consideration. The complexity of recommendations to achieve the goals, the degree of interactions among them, and the relative importance of goals should be considered when deciding how to encode a multi-goal guideline. In general, if there are complex recommendations for each goal, we suggest using a subguideline or an independent guideline to encode them for each goal. If there is a high degree of interactions among recommendations, encoding them in a single clinical algorithm may be necessary. If there is one primary goal and several subsidiary goals, recommendations related to subsidiary goals may be represented as subguidelines.

In the MedSafe-CDS project, we encode the guidelines for HTN, T2DM, and CKD as separate knowledge bases that share a common clinical concept library. Uncontrolled HTN or T2DM—two of the risk factors for CKD—automatically invokes the clinical algorithm for evaluating CKD. A common definition of CKD, shared across the knowledge bases for the three conditions through a library, ensures conceptual consistency and assists with medication recommendations to prevent CKD progression.

Recognizing patient scenarios

Given the episodic nature of chronic disease management, a critical task in applying CIG is the recognition of the applicable patient states at a particular encounter. Current treatment decisions are dependent on past and present assessments of patient status, past decisions and actions, and subsequent outcomes. The challenge is to structure the possible patient states at a high level unambiguously for decision-making at a particular encounter, rather than encoding all of the contextual preconditions of a recommendation separately, with possible duplications. Given the design goal of minimizing clinician data entry, a CDSS should rely on existing EHR data to identify these high-level patient states. The goal of this structural pattern is to give an overarching organization of the clinical algorithm of the encoded guideline to allow both the clinician and the CDSS to focus on the subset of guideline recommendations that are most relevant to the patient.

In the MedSafe-CDS project, we follow the best practice identified by Johnson et al.¹¹ and partition the eligible patient population into disjoint *scenarios*. The first scenario can be either the out-of-scope cases or those missing prerequisite data elements, depending on which group should get messages first. The remaining scenarios are characterized by the states of pharmacological treatment—non-pharmacological treatment, monotherapy, dual therapy, and patients receiving three or more medications—for the signature condition. Treatment decisions then depend on scenarios. For example, if the patient is in the non-pharmacological treatment scenario, the focus of decision making will be whether prescribing a medication is warranted, and if so, which one. In other scenarios, issues such as identifying possible adverse reactions and whether to intensify medication treatment would be more important.

Assessing the patient's current treatment

One aspect of MedSafe-CDS's guideline application is the evaluation of current treatment (primarily medications) and the generation of messages about them. The focus is on current treatment in need of review due to adverse effects, poor control, lack of guideline-concordance, or recently added morbidities. Some consideration of current treatment may be built into clinical algorithms (e.g., replacing contraindicated drugs). Thus, the first step is to build a clinical algorithm that includes all attributes of current treatment that may prompt changes in treatment. The next step is to decide what additional messages, unrelated to treatment change recommendations, are necessary.

The MedSafe-CDS implementation of this aspect of CIG uses the clinical algorithm for treatment-changing attributes of current treatment and uses scenario-specific subguidelines to generate additional comments about the current treatment.

Determining degree of detail in therapeutic recommendations

For some diseases, medications are categorized as first-, second-, or third-line drugs that are appropriate for specific patient populations. In that case, the CDSS may display to clinicians only the “best” choices and hide those that may be less optimal. Alternatively, sometimes it is appropriate to generate for the user a collection of all recommended drug treatment choices, ranked by their properties (e.g., indications and contraindications) as they are applied to the patient case. In this case, the CDSS should be able to give an explanation about the pros and cons for each choice.

Example: A patient with uncontrolled T2DM is currently on non-pharmacological therapy. A guideline may recommend the addition of a first-line drug (metformin) if it is not contraindicated. However, if the first-line drug is contraindicated, then a second-line drug would be recommended if it is not contraindicated. If both first- and second-line drugs are contraindicated, then point the user to additional treatment resources. A CDSS may implement this sequence conservatively, displaying to the user the first-line treatment if the EHR data indicates that it is not contraindicated, or it may display both non-contraindicated first- and second-line drugs all at once. The choice may depend on the goal of CDS, the target audience, or the number of drugs involved. In the case of MedSafe-CDS, we worked with clinicians to determine what they find more helpful: to display a longer or shorter list of recommended drugs stratified by first, second, and third line as per guideline. Our goal is to provide them with an appropriate list of drugs, from which to choose that is not unduly restricted. In the case of T2DM when only metformin is listed as a first line drug, displaying all second-line drug recommendations may be desirable in case a metformin contraindication is found during the current clinic visit which was not present in the EHR at the time the recommendation was generated.

3. Drug information representation

The use of medications is one of the primary therapeutic interventions in managing chronic diseases. A guideline gives detailed information about the usages of, preferences over, and constraints on numerous medications at the drug class and at the individual drug levels. Such information can be further augmented by drug information resources like pharmacy drug use guidance and FDA's Structured Product Label. In this section we discuss ways to represent and use three aspects of medication management.

Representation of drug usage information

As described in a previous publication,²⁰ we can formulate much of drug usage information in a declarative representation format that makes statements about concepts and relationships with no flow-of-control or behavioral assumptions. These statements can be authored and maintained by clinician informaticians with minimal training in the modeling tool and can be used to present compact explanations for recommending these drugs in the CDSS. Declarative representation makes encoding more modular and maintainable and reduces the complexity of clinical algorithms that capture the procedural aspects of guideline recommendations. Figure 1 shows a screenshot of the drug usage information of the medication class thiazide diuretics for the management of hypertension. The drug usage properties include the use of thiazide diuretics as first- or second-line drugs for specific patient populations and its indication and contraindication conditions. They include the formulary-preferred drugs at VA and messages that should be generated for different usages of the drug. Figure 2 shows the information associated with a specific thiazide

diuretic, chlorthalidone. It is a possible drug of choice if the class of thiazide diuretics is being recommended. It shows the maximum recommended daily dose, the ceiling above which the system should not suggest dosage increases, and the fact that it's a preferred drug among thiazide diuretics.

INSTANCE EDITOR
For Instance: ♦ Thiazide Diuretic VA 2020 (instance of Drug_Usage, internal name is ATHENA_HTN_Dashboard_Class1502...)

Label
Thiazide Diuretic VA 2020

Drug Class Name
Thiazide diuretics

Formulary Preferred Drug In Class
♦ chlorthalidone
♦ indapamide
♦ HCTZ

Collateral Actions
♦ Recommend_Add: {relative contra gout uncomplicated, relative contra
♦ Blocked_Add: {Blocked add Thiazide missing K, Blocked add thiazide
♦ Recommend_Add: {General Info Thiazide BPH, Inquire about history o

Absolute Contraindications
• Na ≤ 130 past month
• Gout, complicated
• K ≤ 3 past month & Rx thiazide
• K < 3.5 past month and no Rx thiazide

Compelling Indications

Relative Contraindications
• Gout
• Hyposmality_Hyponatremia
• 130 < Na < 135 past month
• uric acid > 6.8 past year

Relative Indications
• Osteoporosis
• Age >= 65 years

Drug Partners

Drug Partners To Avoid
• Loop Diuretics

Is First-line Drug For
• HTN without CKD with albuminuria (absence of CKD with albuminuria)

Is Second-line Drug For
• HTN and CKD with albuminuria (CKD with albuminuria dx presumed fi

Figure 1. The declarative drug usage information for Thiazide Diuretics as extracted from the VA/DoD 2020 hypertension and 2019 CKD guidelines. This instance of Drug Usage is linked to the Thiazide diuretics class in the medical concept hierarchy.

chlordhalidone (instance of Guideline_Drug, internal name is ATHENA_HTN_Dashboard_...)

Label
chlordhalidone

Drug Usage
♦ Thiazide Diuretic VA 2020

References

Rule In

Rule Out

Being Preferred
♦ true

Generic Drug
• chlorthalidone

Dose Strength Unit

Collateral Acti
♦ Blocked_Increase: {Blocked add
♦ Recommend_Increase: {relative
♦ Recommend_Increase: Inquire

Max Recommended Dose
25.1

Increase Dose Ceiling
24.9

Figure 2. The representation of information associated with the specific guideline drug chlorthalidone: its associated Drug Usage class (Thiazide Diuretic VA 2020), its being a preferred drug (together with indapamide) among thiazide diuretics, maximum recommended drug dose, and dosage ceiling above which no increase-dose recommendation should be made.

Figure 3 shows how such declarative drug information may be used as part of a CDSS drug recommendation. For a patient with diabetes, the system suggests as possible changes in drug regimen, that the dosage of the current medication metformin (1st-line drug) can be increased and that a 2nd-line drug, such as empagliflozin or alogliptin, may be added. A current GFR should be obtained before prescribing alogliptin.

Diabetes Recommendations

II Missing Data

Consider one or more of the following drug options: (Do not change meds if meds recently changed):

☐ View additional details

Increase dose metformin extended release

1st line for Type 2 Diabetes and CKD stages 1-3

Add empagliflozin

2nd line for Type 2 Diabetes and CKD stages 1-3

SGLT2 inhibitors as add-on therapy for T2DM in patients with stage 1 to 3 CKD reduce CKD progression and the risk of cardiovascular events.

Add alogliptin

2nd line for Type 2 Diabetes

II Out-of-date: GFR: older than 1 month 54.0 (7/20/2022).

Figure 3. Screenshot showing a portion of MedSafe-CDS drug recommendation, populated with synthetic data.

Preference ordering of drug classes and specific drugs

A major function of CPGs is to provide guidance for the selection of medications to treat the CPGs' signature diseases. The guidance may be based on effects of the drug class or specific drugs on particular subpopulations of patients that have certain conditions, such as possible contraindications and side effects, interactions with other medications, and effectiveness in reaching therapeutic goals. Some institutions may have preferred drugs, and patients may have preferences for certain drugs (based on attributes of drugs).

As described earlier, we define properties of a Drug Usage class to characterize how the conditions (e.g., congestive heart failure) that define subpopulations' relation (e.g., a relative contraindication) to the use of a drug (e.g., saxagliptin) for the goal of the guideline (e.g., reduce HbA1C). The EON guideline model allows the specification of which properties should be used as the basis of making recommendations. The action node of a clinical algorithm may specify a set of Drug Usage instances that should be evaluated for addition and recommended if, for example, the conditions for which the drugs are the first-line drugs hold (e.g., metformin, being the first-line drug for patients with diabetes). A scoring function is used to translate the multiple conditions as they apply to a particular patient into a preference ordering (e.g., a drug with more compelling indications for a patient gets a higher score).

Within a drug class that has a common set of usage criteria (e.g., indications/contraindications, good/bad drug partners), we need to be able to express choice and preference among the specific drugs with a drug class. The EON guideline model allows the specification of a set of "formulary-preferred drugs" as members of a Drug Usage instance. In Guideline Drug representation shown in Figure 2, we use rule-in/rule-out criteria to come up with a set of drug candidates among the formulary-preferred drugs. Among the drug candidates, there is the "being preferred condition" to determine whether a drug is preferred among the candidates.

Constraint over drug selection

Above and beyond detailed drug usage information and preferences among possible drug choices, there may be constraints on drug usage. In this section we will discuss one common constraint: drug classes for which at most one drug of the same class should be given. This constraint may arise because it is not desirable to prescribe more than one of the drugs that have a similar mechanism of action. In that case, the CDSS needs to stop recommending adding a drug since another drug of the same class is already on the patient's medication list. Here the problem is how "drugs of the same class" are represented in the system and how to specify which drug class has this property, given that a specific drug may be classified in multiple ways. One possibility is to have a designated drug class with subclasses in a drug hierarchy. Another possibility is to have an enumeration of drugs that should not be given together.

In a MedSafe-CDS guideline encoding, there are three ways of handling the "At most one drug of the same class" problem:

1. If the drug class has a representation as a Drug Usage instance (e.g., ACE inhibitor), and the MedSafe CDS KB evaluates whether to add the Drug Usage instance, the execution engine will automatically rule out all members of the drug class if a member of the class is already in the patient's medication list.

2. If the drug class does not have a representation as a single Drug Usage instance, then one way to handle this is to create a query for drugs of the class and add it to “do-not-start-conditions” property of the Drug Usage instances corresponding to members of the drug class.
3. Alternatively, without creating a query, the specific drugs to avoid can be added to the “drug partners to avoid” property of the Drug Usage instances corresponding to members of the drug class.

Having a single Drug Usage instance for all members of the drug class is the cleanest way to handle the problem. Failing that, the “drug partners to avoid” method may be most tedious to maintain, as one must enter the members of the drug class in all relevant instances of Drug Usage. However, having such a representation allows this information to be used in query expressions elsewhere.

Discussion

Patel et al.²⁴ argued that the limited acceptance of CPGs may be because of their semantic complexity and that improved understanding of the semantics and structure of CPGs may help to improve the match between a guideline’s recommended actions and the physician-user’s mental models of what seems appropriate in a given case. The enumeration of the structural patterns in chronic disease guidelines makes explicit the high-level mental models clinicians may have for the management of patients with chronic diseases, and links them to knowledge representation and CDS possibilities as exemplified by MedSafe-CDSS.

A theme in the description of the structural patterns and their use in CDS is the complexity of medical knowledge and decision making in chronic disease management. Drug recommendations involve the evaluation of several possible drugs along multiple dimensions of drug properties and clinical characteristics. It is difficult to encapsulate all of these considerations in deterministic flowcharts or simple reminders.

A limitation of the MedSafe-CDS approach is the lack of explicit considerations of patient preferences. By providing multiple treatment alternatives with explanations for each, we may facilitate clinicians/patients shared decision making. Future work could incorporate the elicitation of patient goals²⁵ and preferences into shared decision making for clinical management.²⁶

Another limitation of the structural patterns reported in this paper is the difficulty in formally evaluating them. These patterns emerged over years of iteratively encoding different versions of multiple chronic disease guidelines. While we expect the issues raised in these patterns (e.g., data availability, multiple goals, drug usage and constraints) are common in encoding chronic disease guidelines, the particular MedSafe-CDSS solutions are necessarily specific to its guideline representation and are illustrative in nature. In one sense they were validated when the same patterns were applied to a new guideline and found useful during MedSafe-CDSS development. In another sense they were validated when the MedSafe-CDSS knowledge bases, encoded using these patterns, were found to produce correct outputs for multiple test cases.^{16,27} A true evaluation, involving the use of these patterns in a completely different CDSS and thus verifying the applicability of the patterns in new settings, is beyond the scope of a single project.

The structural patterns reported in this paper focus on the management of chronic diseases according to CPGs. Chronic disease guidelines tend to present clinical algorithms for screening, diagnosis, monitoring, and management of one disease. However, chronic conditions often exacerbate and multiply; consequently, chronic care involves coordination across multiple clinical domains and often multiple guidelines may be relevant for the management of a single patient²⁸. A significant literature has emerged on the problems of detecting and mitigating interactions arising from the use of multiple CPGs for CDS.^{29, 30}

MedSafe-CDSS does not have formal methods to detect and mitigate interactions among CPG recommendations. Instead, our approach relies on the incorporation of comorbidities considerations within a CPG document³¹ and careful knowledge synthesis of recommendations from multiple guidelines to avoid conflicting recommendations. Our declarative representation of the indications, contraindications, and other drug properties relevant to the use of drugs and drug classes facilitated the development of a measure of comorbidity interrelatedness (the degree to which patients’ co-occurring conditions interact to generate clinical complexity).³²

The American Geriatrics Society (AGS) Expert Panel on the Care of Older Adults with Multimorbidity recommends a management approach that will consider the multiple problems particular to each individual with multimorbidity.³³ The approach goes beyond evidence-based CPGs to combine patient preferences; interpretation of the evidence; framing of decision in the context of risks, burdens, benefits, and prognosis; clinical feasibility; and optimization of therapies and care plan. It is still a challenge for CDSSs to incorporate the more holistic framing of chronic disease management as recommended by the AGS expert panel. Potential use of CDS to address medical management choices

should also consider how CDS can be embedded in systems of care that address behavioral factors in health care, as well as services to address social determinants of health.

Conclusion

We abstracted nine structural patterns in the modeling of chronic disease CPGs for the purpose of providing clinical decision support. These patterns guide the encoding of the CPG recommendations, resulting in the ability of MedSafe-CDSS to provide focused yet complex decision support.

Acknowledgement

This work was supported in part by Quality Enhancement Research Initiative (QUERI) project: Optimizing Appropriate Use of Medications for Veterans, IP1 HX002261-01 (PIs: Heidenreich, Goldstein) from the United States (U.S.) Department of Veterans Affairs Health Services Research and Development. Views expressed are those of the authors and not necessarily of the Department of Veterans Affairs or the United States government.

References

1. Peleg M. Computer-interpretable clinical guidelines: a methodological review. *Journal of biomedical informatics* 2013; 46: 744-763. Review. DOI: 10.1016/j.jbi.2013.06.009.
2. Shiffman RN, Michel G, Essaihi A, et al. Bridging the Guideline Implementation Gap: A Systematic, Document-Centered Approach to Guideline Implementation. *J Am Med Inform Assoc* 2004; 11: 418-426.
3. Boxwala AA, Rocha BH, Maviglia S, et al. A multi-layered framework for disseminating knowledge for computer-based decision support. *J Am Med Inform Assoc* 2011; 18 Suppl 1: i132-139. 20111103. DOI: 10.1136/amiajnl-2011-000334.
4. Peleg M, Boxwala AA, Tu SW, et al. The InterMed approach to sharable computer-interpretable guidelines: a review. *Journal of the American Medical Informatics Association* 2004; 11: 1-10. DOI: 10.1197/jamia.M1399.
5. HL7. FHIR Clinical Guidelines (v1.0.0), <http://hl7.org/fhir/uv/cpg/index.html> (2022, accessed 2022/02/10 2022).
6. Peleg M and Tu SW. Design patterns for clinical guidelines. *Artif Intell Med* 2009; 47: 1-24. 2009/06/09. DOI: S0933-3657(09)00083-9 [pii]10.1016/j.artmed.2009.05.004.
7. Kaiser K and Marcos M. Leveraging workflow control patterns in the domain of clinical practice guidelines. *BMC Med Inform Decis Mak* 2016; 16: 20. 20160210. DOI: 10.1186/s12911-016-0253-z.
8. Moser M and Miksch S. Improving Clinical Guideline Implementation through Prototypical Design Patterns. In: *10th Conference on Artificial Intelligence in Medicine (AIME 05)* Aberdeen, Scotland, 2005.
9. Gooch P and Roudsari A. Computerization of workflows, guidelines, and care pathways: a review of implementation challenges for process-oriented health information systems. *J Am Med Inform Assoc* 2011; 18: 738-748. 20110701. DOI: 10.1136/amiajnl-2010-000033.
10. Gamma E, Helm R, Johnson R, et al. *Design Patterns: Element of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
11. Johnson PD, Tu S, Booth N, et al. Using scenarios in chronic disease management guidelines for primary care. *Proc AMIA Symp* 2000: 389-393. 2000/11/18. DOI: D200629 [pii].
12. Tinetti ME, Fried TR and Boyd CM. Designing health care for the most common chronic condition--multimorbidity. *JAMA* 2012; 307: 2493-2494. 2012/07/17. DOI: 10.1001/jama.2012.5265.
13. Goldstein MK, Hoffman BB, Coleman RW, et al. Implementing clinical practice guidelines while taking account of changing evidence: ATHENA DSS, an easily modifiable decision-support system for managing hypertension in primary care. *Proc AMIA Symp* 2000: 300-304. 2000/11/18. DOI: D200559 [pii].
14. Goldstein MK, Coleman RW, Tu SW, et al. Translating research into practice: organizational issues in implementing automated decision support for hypertension in three medical centers. *J Am Med Inform Assoc* 2004; 11: 368-376. DOI: 10.1197/jamia.M1534.
15. Tu SW, Martins S, Oshiro C, et al. Automating Performance Measures and Clinical Practice Guidelines: Differences and Complementarities. *AMIA Annu Symp Proc* 2016; 2016: 1199-1208.
16. Usman OA, Oshiro C, Chambers JG, et al. Selecting Test Cases from the Electronic Health Record for Software Testing of Knowledge-Based Clinical Decision Support Systems. *AMIA Annu Symp Proc* 2018; 2018: 1046-1055. 2019/04/26.
17. Musen MA. The Protégé project: A look back and a look forward. *AI Matters* 2015; 1. DOI: 10.1145/2557001.25757003.
18. Tu SW and Musen MA. Modeling data and knowledge in the EON guideline architecture. *Stud Health Technol Inform* 2001; 84: 280-284. 2001/10/18.

19. Tu SW and Musen MA. Representation Formalisms and Computational Methods for Modeling Guideline-Based Patient Care. In: Heller B, Loffler M and Musen MA (eds) *Computer-based Support for Clinical Guidelines and Protocols*. Amsterdam, The Netherlands: IOS Press, 2001, pp.125-142.
20. Tu SW, Hrabak KM, Campbell JR, et al. Use of declarative statements in creating and maintaining computer-interpretable knowledge bases for guideline-based care. *AMIA Annu Symp Proc* 2006; 784-788.
21. Tu SW and Musen MA. From guideline modeling to guideline execution: defining guideline-based decision-support services. *Proc AMIA Symp* 2000; 863-867. 2000/11/18. DOI: D200735 [pii].
22. Tso GJ, Tu SW, Oshiro C, et al. Automating Guidelines for Clinical Decision Support: Knowledge Engineering and Implementation. *AMIA Annual Symposium proceedings / AMIA Symposium AMIA Symposium* 2016; 2016: 1189-1198.
23. Goldstein MK, Hoffman BB, Coleman RW, et al. Implementing clinical practice guidelines while taking account of changing evidence: ATHENA DSS, an easily modifiable decision-support system for managing hypertension in primary care. *Proceedings / AMIA Annual Symposium AMIA Symposium* 2000; 300-304.
24. Patel VL, Arocha JF, Diermeier M, et al. Methods of cognitive analysis to support the design and evaluation of biomedical systems: the case of clinical practice guidelines. *J Biomed Inform* 2001; 34: 52-66. DOI: 10.1006/jbin.2001.1002.
25. Hsu KY, Slightam C, Shaw JG, et al. High-Need Patients' Goals and Goal Progress in a Veterans Affairs Intensive Outpatient Care Program. *J Gen Intern Med* 2019; 34: 1564-1570. 20190528. DOI: 10.1007/s11606-019-05010-w.
26. Das AK, Ahmed BA, Garten Y, et al. Knowledge-based method for building patient decision-analytic tools. *AMIA Annu Symp Proc* 2006; 175-179. 2007/01/24. DOI: 85650 [pii].
27. Tso GJ, Yuen K, Martins S, et al. Test Case Selection in Pre-Deployment Testing of Complex Clinical Decision Support Systems. *AMIA Joint Summits on Translational Science proceedings AMIA Summit on Translational Science* 2016; 2016: 240-249.
28. Boyd CM, Darer J, Boult C, et al. Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases: implications for pay for performance. *Jama* 2005; 294: 716-724.
29. Riaño D and Ortega W. Computer technologies to integrate medical treatments to manage multimorbidity. *J Biomed Inform* 2017; 75: 1-13. 20170921. DOI: 10.1016/j.jbi.2017.09.009.
30. Bilici E, Despotou G and Arvanitis TN. The use of computer-interpretable clinical guidelines to manage care complexities of patients with multimorbid conditions: A review. *Digit Health* 2018; 4: 2055207618804927. 20181003. DOI: 10.1177/2055207618804927.
31. Leung TI, Jalal H, Zulman DM, et al. Automating Identification of Multiple Chronic Conditions in Clinical Practice Guidelines. *AMIA Joint Summits on Translational Science proceedings AMIA Summit on Translational Science* 2015; 2015: 456-460. 20150325.
32. Zulman DM, Martins SB, Liu Y, et al. Using a Clinical Knowledge Base to Assess Comorbidity Interrelatedness Among Patients with Multiple Chronic Conditions. *AMIA Annu Symp Proc* 2015; 2015: 1381-1389. 20151105.
33. Multimorbidity AGSEPOTCoOA. Guiding principles for the care of older adults with multimorbidity: an approach for clinicians: American Geriatrics Society Expert Panel on the Care of Older Adults with Multimorbidity. *J Am Geriatr Soc* 2012; 60: E1-e25. 20120919. DOI: 10.1111/j.1532-5415.2012.04188.x.