Hierarchical representation of differential diagnosis lists for clinical decision support systems

Takashi Okumura
National Institute of Public Health, Japan
2-3-6 Minami, Wako city
Saitama Pref., Japan
taka@niph.go.jp

Masaki Tagawa
Kyushu Institute of Technology
680-4 Kawazu, Iizuka city
Fukuoka Pref., Japan
masaki@club.kyutech.ac.jp

ABSTRACT

Clinical decision support systems (CDSSs) can generate differential diagnosis lists that may contain hundreds of diseases. These lists grow in size as coverage expands to rare diseases, but large lists can easily become a burden on user cognition. To address this issue, we first outline the representations of differential diagnosis lists on current CDSSs, and then propose a novel approach that represents these differential diagnosis lists hierarchically, coupled with an algorithm for optimal initialization. Preliminary evaluation suggested that our proposed approach outperforms existing approaches with respect to search costs, particularly for large lists. This hierarchical representation should alleviate the cognitive load on user physicians and provide an efficient means to search through very large lists.

Categories and Subject Descriptors

H.5.2 [**Information Systems**]: Information interfaces and presentation

General Terms

Clinical Decision Support Systems, Expert systems

Keywords

Differential diagnosis, Hierarchical representation

1. INTRODUCTION

A differential diagnosis list is a list of possible diagnosis that is tentatively built for planning further examinations aimed toward an accurate diagnosis. Physicians are educated to consider at least three possible diagnoses in Japan: *common*, *curable*, and *critical* diseases, and these three can be a minimal set for differential diagnosis. For most situations, however, there are several more possibilities to be considered. Nevertheless, a list should not exceed 20 or 30 items, because this exceeds the capacity of human cognition. In these cases, physicians set up a diagnostic hypothesis to focus on a more limited number of possibilities.

IDPHA 2014, May 20-22

DOI 10.4108/icst.pervasivehealth.2014.255785

In contrast, clinical decision support systems (CDSS) can generate a differential diagnosis list with a far greater volume. This list is built on a set of given findings, and provides a diagnostic probability for each item. The resulting list is an unstructured flat list, sorted by the probabilities, although such a list becomes too monotonous to browse as the number of listed items grows. This situation occurs during the diagnosis of rare diseases, if physicians have no choice but to use general search engines by providing key findings of a case as their search terms [10]. In these cases, physicians are forced to laboriously search through a large list of either disease names, or web pages.

To develop a CDSS that covers various rare diseases, improving the representation would be indispensable. Accordingly, this article addresses this problem, through a hierarchical representation of a differential diagnosis list. This structured representation should alleviate the cognitive load placed on user physicians and provide an efficient means to formulate and test their clinical hypotheses. This should also contribute to the usability of related health-care applications.

The sections in this article are organized as follows. Section 2 provides a survey of existing representations on relevant CDSSs. Section 3 provides the hierarchical representations of differential diagnosis lists, and an algorithm for efficient initialization. Section 4 demonstrates the results of our preliminary evaluation. Section 5 discusses the proposed representation, and Section 6 concludes this paper.

2. BACKGROUND AND RELATED WORK

Although there are a various types of CDSSs [2], which are not limited to those that output differential diagnosis lists, they form the mainstream of research and production systems. Many systems in this category generate flat lists of possible diagnoses that are sorted by their possibilities (Figure 1-a), using given sets of clinical findings as their inputs [11, 12]. Such a simple listing is a good fit, particularly for people without clinical knowledge [12], as the list only includes a limited set of common diseases and this simplicity is well suited for such users. Even experts may prefer a flat listing if the context is limited, e.g., the diagnosis of infectious diseases only [11]. However, for hard-to-diagnose cases, the situation changes, because a CDSS may generate hundreds of possible diagnoses. Such a list can easily overwhelm human cognition if the candidate diseases are presented in a simple list that requires a serial search.

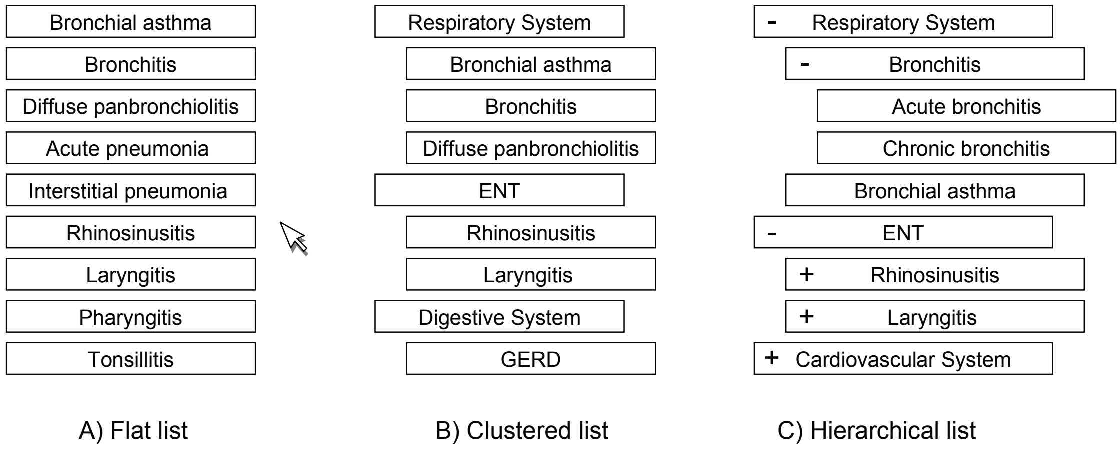

Figure 1: Representation of Differential Diagnosis list

To aid physicians in finding "a needle in a haystack", a common approach is to cluster the diseases into groups, such as organ systems to aggregate similar diseases (Figure 1-b). This is a standard style for a modern CDSS [8, 6, 1]. Nevertheless, for a large list, a single cluster can contain a considerable number of items, which again can overwhelm users. For example, there are hundreds of autoimmune diseases and a homogeneous cluster would certainly become a burden on user cognition.

As illustrated, CDSSs have mostly used flat listing or clustered listing for their output. A few exceptions include paneling of pictures for visual diagnosis in dermatology [9], but this also necessitates serial searching. These representations increase the cognitive load, as system coverage expands to rare diseases. To address this inefficiency in representation, a natural approach would be the recursion of clustering (Figure 1-c). Although the hierarchical representation of diseases appears to be reasonable, such an approach has not been documented yet, to the best of our knowledge.

3. HIERARCHICAL REPRESENTATION OF DIFFERENTIAL DIAGNOSIS LIST

3.1 Clinical requirements and Technical interpretation

The last section argued that listings with a vast number of diseases cognitively burden user physicians and that a hierarchical representation may improve search efficiency. However, various questions remain for how to formulate and present this hierarchy. For example, the preferred depth of the hierarchy is unknown and there is no established way to cluster diseases. In a particular representation, it is also unknown if the nodes should be collapsed or expanded in the initial state and how they are related to clinical concepts.

To address these issues, the requirements from the clinical perspective can be summarized as follows. First, disease clustering is preferable when the number of items involved increases. Second, clustering contributes to lowering the cognitive load, particularly for homogeneous clusters. Third, in contrast, a detailed listing is preferable if there is diversity in the listed items. Fourth, clustering must be related to clinical concepts, to help physicians navigate in the searching process based on their clinical knowledge. Finally, there must be a certain limit for the number of diseases presented at a time, not to overwhelm the users.

These clinical requirements must be technically interpreted, to implement an actual CDSS, and summarized as follows. First, a hierarchical representation for a differential diagnosis should provide details for more likely diseases but necessary aggregation is required for less likely diseases. Second, the hierarchy must be formulated with clinical perspective. Third, not to overload the user physicians, the number of visible diseases should not exceed a certain limit.

3.2 An unfolding algorithm

To meet these goals, a hierarchical tree must be initialized as such. There are two primary ways for such an initialization: either collapse a hierarchical tree of diseases until the resulting tree meets the requirements or unfold a collapsed tree, unless a particular state violates the required goals. Because the number limit ($<$ 20 or 30) is much smaller than the number of diseases in the tree, which can be several hundreds of diseases, a collapse algorithm would be inefficient. Consequently, this article presents an unfolding algorithm that initializes the tree into an optimal state using finite steps.

Figure 2 illustrates our algorithm in a pictorial way. First, the algorithm selects *node 0* (1-1). If the number of items in *node 0* satisfies the limit, then the node unfolds (1-2). During the second round, the algorithm selects nodes on the *frontier*, which is the node to inspect for unfolding, and defined as the topmost item at each depth in the tree. In this case, *node 1-1* is on the frontier (2-1). If adding items does not violate the limit, then *node 1-1* unfolds (2-2). Third, *node 2-1* (A) and *node 1-2* (B) are on the frontier (3-1), and these nodes are evaluated to determine which node to unfold. For this purpose, the variances of the probability scores for the child nodes are calculated and used to sort the frontier nodes. Then, the top node is checked to see if the unfolding violates the predefined number limit of visible items (3-2). The succeeding rounds serve in a similar manner, and continue until the number limit is reached, or until the frontier disappears.

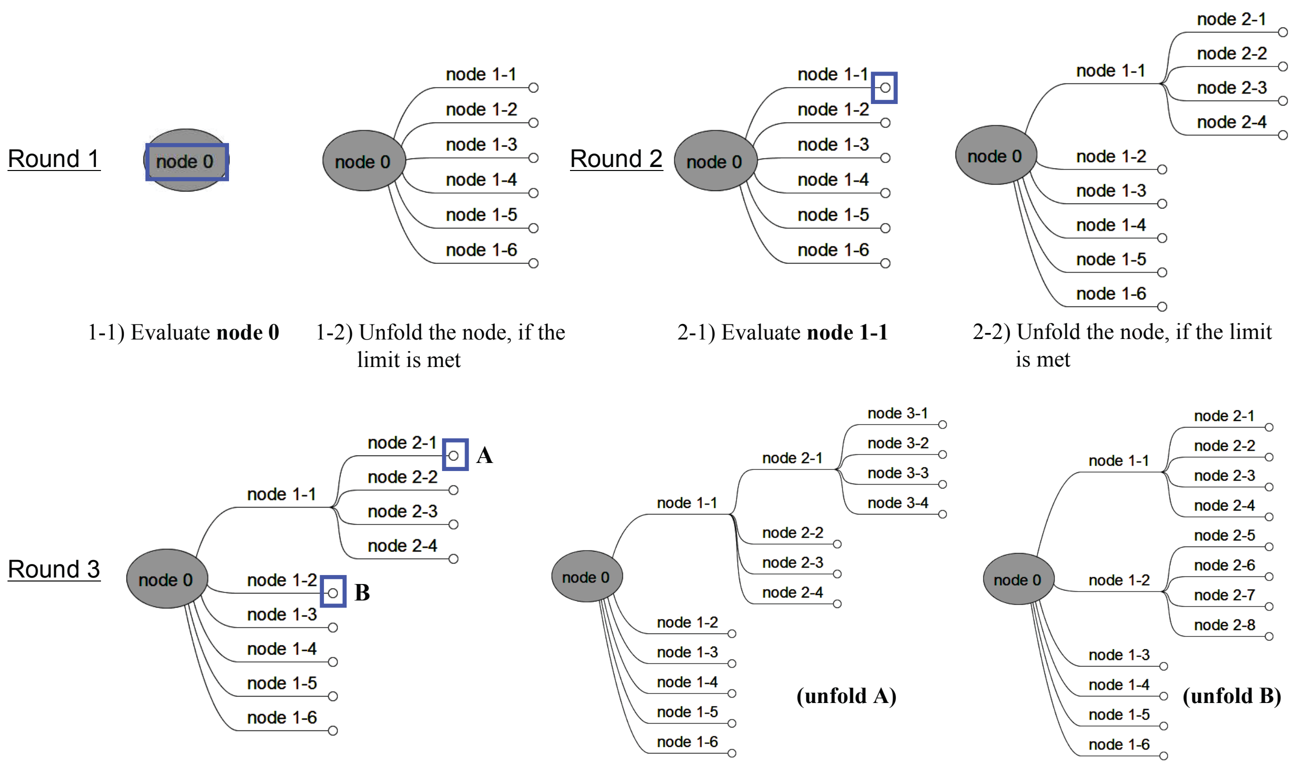

1-1) Evaluate **node 0** 1-2) Unfold the node, if the limit is met 2-1) Evaluate **node 1-1** 2-2) Unfold the node, if the limit is met

3-1) Evaluate **node 2-1** and **node 1-2** 3-2) Unfold the one on the top, if the limit is met, and repeat the process

Figure 2: Unfolding algorithm for optimal initialization

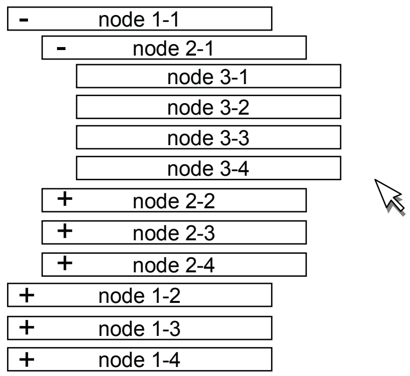

Figure 3: Graphical representation of tree 3-2 (A)

This algorithm transforms a differential diagnosis list into a hierarchical structure with a desirable initial state for each node. An example is shown in Figure 3, which depicts state 3-2 (A) in Figure 2 in a possible user interface style. Starting from the initial state, users can select a node to focus upon, and then expand this node to look into the detail, to search for a more reasonable diagnosis. The search cost depends on the structure of the tree and the number of visible nodes. The next section investigates the cost involved, and demonstrates the advantage of this proposed representation by comparisons to existing approaches.

4. EVALUATION

Estimating the search costs for each representation requires realistic assumptions. In particular, a hierarchical representation requires a systematic classification model of diseases; for this purpose, ICD-10 [13] may provide a basis for hierarchical organization of diseases in a differential diagnosis list. This classification system comprises 22 chapters, each of which contains items for second, third and fourth level domains. In this tree, the second level contains 263 *blocks* (12 blocks per chapter), the third level contains 2048 *codes* (8 codes per block), and the fourth level contains 10083 diseases (5 diseases are mapped to one code, on average).

For this evaluation, the search costs for the representations in Figure 1 were compared by changing the number of items in a list. A flat list simply grows, in response to adding items to the list. A clustered list requires a parameter for how to group the items, and for this purpose, the number of ICD code per block is used. Finally, a hierarchical representation reused the tree parameters of ICD-10. In this evaluation, the number of visible nodes was set to 10, for the initialization of the hierarchical representation. Note that this limit was not applied to the flat and the clustered lists, because the limitation makes these two approaches incomparable to a hierarchical list.

Figure 4 illustrates the search costs for these representations. Their average costs are shown in lines and the worst case costs are in positive error bars. As illustrated, the flat list shows linear increase in the cost, exhibiting inefficiency

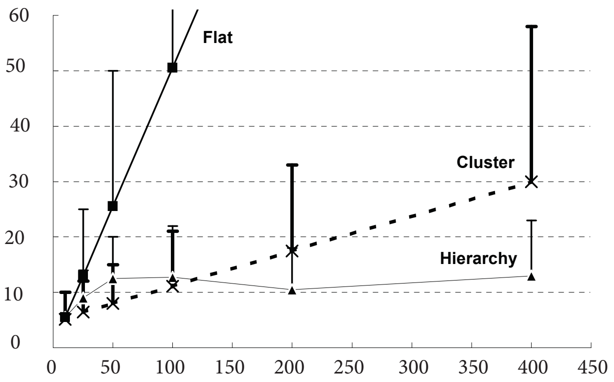

Figure 4: Search cost against # of items

in the search. The cluster approach also shows the linearity, because the unit size of the cluster is fixed in the evaluation. In contrast, there was a logarithmic boundary for the hierarchical representations, and it outperformed the clustered list in large list settings. The worst case performances also suggested advantage of the hierarchical approach.

5. DISCUSSION

For the diagnoses of rare diseases and hard-to-diagnose cases, differential diagnosis lists can become large and physicians are expected to consider even those items at the bottom of the differential diagnosis list [5]. Accordingly, it is reasonable for a CDSS to provide a reasonable method to navigate through a very large list, in addition to improving the accuracy of the listings.

To efficiently display such a large list, Human-Computer Interaction (HCI) studies have used hierarchical representation, as a standard approach. For example, attempts have been made for hierarchical representation of unstructured video streams [7], which can be applied to medical field [14]. There is also a model for hierarchical representation of topics, a so-called "degree of interest tree" [3], but this has not been applied to presentations of differential diagnosis lists.

In this article, we applied an orthodox technique to ease the cognitive load imposed by a differential diagnosis list and proposed a simple algorithm to adjust the initial state of the representation from a clinical perspective. The proposed representation should accelerate the searching process by providing means to prune the search tree based on clinical reasoning and should allow physicians to efficiently identify diseases of interest, which can be lost among similarity and complexity of diseases.

For evaluation of CDSSs, the probability of including a correct diagnosis in the topmost part of the list is sometimes used [4]. This strategy appears reasonable, as a serial search of an entire list is costly. On the other hand, a hierarchical representation can compress a large listing into a compact form and virtually broaden the scope for users, which essentially has the same effect as improving the accuracy of diagnostic algorithm. Because the latter involves various difficulties, particularly for diagnosing rare diseases, a hierarchical representation may be a practical solution.

6. CONCLUSIONS

Our proposed approach can hierarchically represent a very large differential diagnosis list in an organized manner, which is sufficiently handy to search the entire space. This reasonable presentation originates from the clinical perspective, and materialized with technical and cognitive expertise. Although cognitive and clinical evaluations are left for future work, it is likely that this type of representation will contribute to higher usability and efficiency of modern CDSSs.

7. REFERENCES

[1] G. O. Barnett, J. J. Cimino, J. A. Hupp, and E. P. Hoffer. Dxplain: an evolving diagnostic decision-support system. *Jama*, 258(1):67–74, 1987.

[2] E. S. Berner. *Clinical Decision Support Systems*. Springer, 2007.

[3] S. K. Card and D. Nation. Degree-of-interest trees: A component of an attention-reactive user interface. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 231–245. ACM, 2002.

[4] M. L. Graber and A. Mathew. Performance of a web-based clinical diagnosis support system for internists. *Journal of general internal medicine*, 23(1):37–40, 2008.

[5] M. M. Grinsell and V. F. Norwood. At the bottom of the differential diagnosis list: unusual causes of pediatric hypertension. *Pediatric Nephrology*, 24(11):2137–2146, 2009.

[6] MedTech USA, Inc. DiagnosisPro. http://www.diagnosispro.com/.

[7] D. Ponceleon and A. Dieberger. Hierarchical brushing in a collection of video data. In *System Sciences, 2001. Proceedings of the 34th Annual Hawaii International Conference on*, pages 8–pp. IEEE, 2001.

[8] P. Ramnarayan, A. Tomlinson, A. Rao, M. Coren, A. Winrow, and J. Britto. Isabel: a web-based differential diagnostic aid for paediatrics: results from an initial performance evaluation. *Archives of disease in childhood*, 88(5):408–413, 2003.

[9] I. M. Tleyjeh, H. Nada, and L. M. Baddour. Visualdx: decision-support software for the diagnosis and management of dermatologic disorders. *Clinical infectious diseases*, 43(9):1177–1184, 2006.

[10] Y. Tokuda, M. Aoki, S. B. Kandpal, and L. M. Tierney. Caught in the web: E-diagnosis. *Clinical Care Conundrums: Challenging Diagnoses in Hospital Medicine*, pages 169–175, 2013.

[11] L. Y. Victor and S. C. Edberg. Global infectious diseases and epidemiology network (gideon): a world wide web-based program for diagnosis and informatics in infectious diseases. *Clinical infectious diseases*, 40(1):123–126, 2005.

[12] WebMD, LLC. . Symptom Checker. http://symptoms.webmd.com/.

[13] World Health Organization. *ICD-10: International statistical classification of diseases and related health problems: 10th revision, Version:2010*. World Health Organization, 2010.

[14] X. Zhu, J. Fan, A. K. Elmagarmid, and W. G. Aref. Hierarchical video summarization for medical data. In *Electronic Imaging 2002*, pages 395–406. International Society for Optics and Photonics, 2001.