

Evaluation of Hierarchical Clustering via Markov Decision Processes for Efficient Navigation and Search

Raul Moreno, Wěipéng Huáng, Arjumand Younus, Michael O’Mahony, and Neil J. Hurley

Insight Centre for Data Analytics, University College Dublin, Ireland
`{firstname.surname}@insight-centre.org`

Abstract. In this paper, we propose a new evaluation measure to assess the quality of a hierarchy in supporting search queries to content collections. The evaluation measure models the scenario of a searcher seeking a particular target item in the hierarchy. It takes into account the structure of the hierarchy by measuring the cognitive challenge of determining the correct path in the hierarchy as well as the reduction in search time afforded by hierarchy. The goal is to propose a general-purpose measure that can be applied in different application contexts, allowing different hierarchical arrangements of content to be quantitatively assessed.

1 Introduction

Content collections are commonly arranged in hierarchical taxonomies. For example, a commercial retailing site in the clothing sector may arrange its catalogue in categories such as “Leisurewear”, which may contain a sub-category “Sportswear”. A customer with a particular product in mind can then search through the category hierarchy in order to zone in on the sub-set of the collection that contains the required target product. Generally, there is more than one way to organise a hierarchy. For example, a book such as “Harry Potter” might be stored under a high-level category “Children’s Books” and sub-category “Fantasy”, or indeed “Fantasy” might be the higher-level category with “Children’s Books” beneath it. Or a different taxonomy might instead use a category of “Books about Wizards”. Given a particular hierarchical organisation of content, we address the question of how to evaluate its usefulness for supporting content search queries. Such a measure would facilitate the quantitative comparison of one hierarchy with another, without the need for a ground-truth hierarchy.

A range of methods for measuring the coherence of sub-categories within a hierarchy have been proposed [1]. However, we seek a measure that evaluates the *structure* of the hierarchy as well as the coherence of the hierarchical clusters. In particular, we seek a measure that addresses the question of how easy it is to navigate through a hierarchy from its root to some target content, accounting for the cognitive cost of choosing a correct path at each branch of the hierarchy. A few hierarchical clustering measures address this issue [2, 3]. Johnson et al.

[3] evaluate hierarchy structure but their approach relies on the availability of ground-truth clusters. Cigarran et al’s approach [2] proposes a goal-oriented evaluation measure of the hierarchical clustering quality, which considers the content of the cluster, the hierarchical arrangement and the navigation cost. Notably, the approach utilizes the idea of a Minimal Browsing Area (MBA)¹ to measure navigation cost. We differ from these approaches in that our focus is on evaluation without ground-truth and, based on Markov decision processes [4, 5] our novel evaluation measure is designed to model the behaviour of a searcher navigating a hierarchy.

2 Markov Decision Processes

Markov decision processes (MDPs) [4, 5] provide an appropriate mathematical framework for modelling decision making where outcomes are partly random and partly under the control of a decision maker. An MDP is a discrete time process, in which at each time step, the process is in some state s and the decision-maker must choose an available action a to move to a new state s' . Each move has an associated reward $R_a(s, s')$, which provides the motivation for choosing particular decisions. The goal of an MDP may be stated in terms of seeking a *policy* for choosing the action at each step that earns the maximum expected cumulative reward over time.

2.1 Navigating a Hierarchy as an MDP

In our scenario, the searcher must make a decision at each visited node in the hierarchy. There are two possible actions at each state (or node in the hierarchy):

1. Stay at the current node and examine the documents stored under it
2. Navigate to a child node

Assuming that the target document is stored under the current node i , the reward of choosing to remain at this node may be written as the overall reduction in the number of documents (n_i) that need to be examined, compared with the full document set (n). Navigating to the child containing the target accumulates reward 0 at the current step, however opens the possibility of an enhanced reward at a subsequent step, when the user finally decides to explore the documents at a lower-level node. Once the target document is examined, the search process ends. If an incorrect child is chosen, then a fixed reward of -1 is obtained regardless of subsequent steps, which can never reach the target. Hence, it is unnecessary to examine the search any further after a bad choice is made. In summary, the search process consists of visits to a set of states $\{s_0, s_1, \dots\}$ containing the target document with the following rewards assigned to each choice:

$$R_{\text{remain}}(s_i) = 1 - \frac{n_i}{n} \quad R_{\text{good}}(s_i, s_{i+1}) = 0 \quad R_{\text{bad}}(s_i) = -1.$$

¹ An MBA for the target is the minimal part of the hierarchy that, starting from the top node, a user must explore in order to reach all relevant items.

Effectively, the search ends once a non-zero reward is obtained.

The expected reward depends on the navigation policy π that the searcher follows in order to decide between the alternatives at each step in the process. Given a particular policy, π , we can evaluate the probabilities associated with the different possible state transitions. In particular, given a target document t , we write the probabilities at node i of the hierarchy as:

- $p_i(t)$ = probability of choosing the correct path to the target
- $q_i(t)$ = probability of choosing an incorrect path
- $r_i(t)$ = probability of choosing to examine all documents under the node i

We have that $p_i(t) + q_i(t) + r_i(t) = 1$. Also, $q_i(t)$ is the sum of the probabilities to navigate to all incorrect nodes at the level, which indicates the tree structure is not necessarily binary.

Let $\{s_0, s_1, \dots, s_{d(t)}\}$ be the set of nodes on the path from the root ($= s_0$) to the leaf node, $s_{d(t)}$, that contains the target t . Finally, we propose to define the *hierarchy score* given target document t as the expected reward, i.e.,

$$\text{HQ}(L, \{t\}, \pi) = \mathbb{E}[R(\pi)] = \sum_{i=1}^{d(t)} \left(\prod_{j=0}^{i-1} p_j(t) r_i(t) \left(1 - \frac{n_i}{n}\right) - \prod_{j=0}^{i-1} p_j(t) q_i(t) \right),$$

where $p_0 = 1$. The value of the hierarchy for searching a document set D using policy π is

$$\text{HQ}(L, \pi) = \sum_{t \in D} w_t \text{HQ}(L, \{t\}, \pi),$$

where w_t is the importance weight of document t such that $\sum_t w_t = 1$ (e.g. w_t may be taken as the relative frequency of a search for t).

As an illustration of the formula, consider a hierarchy in which the number of documents is reduced by a factor of ϵ at each level of the hierarchy. Thus, the reward if the target is found at level i , is $1 - \epsilon^i$. We examine the formula in the limit as $i \rightarrow \infty$, for different choices of the three decision probabilities.

In particular, considering that $p_i = \gamma$, $r_i = \alpha(1 - \gamma)$ and $q_i = (1 - \alpha)(1 - \gamma)$ for some $\alpha \in [0, 1], \gamma \in [0, 1]$, we have

$$\begin{aligned} \mathbb{E}[R] &= \sum_{i=0}^{\infty} \gamma^i \alpha(1 - \gamma)(1 - \epsilon^i) - \sum_{i=0}^{\infty} \gamma^i (1 - \alpha)(1 - \gamma) \\ &= (1 - \gamma) \left(\sum_{i=0}^{\infty} (\alpha - (1 - \alpha)) \gamma^i - \alpha \sum_{i=0}^{\infty} (\gamma \epsilon)^i \right) \\ &= (1 - \gamma) \left((2\alpha - 1) \sum_{i=0}^{\infty} \gamma^i - \alpha \sum_{i=0}^{\infty} (\gamma \epsilon)^i \right) \\ &= (1 - \gamma) \left(\frac{2\alpha - 1}{1 - \gamma} - \frac{\alpha}{1 - \gamma \epsilon} \right) \\ &= (2\alpha - 1) - \alpha \frac{1 - \gamma}{1 - \gamma \epsilon}. \end{aligned}$$

With $\gamma > 0$ and $\alpha \rightarrow 0$, the searcher always descends through the hierarchy, with a non-zero probability of taking the wrong branch at each descent, so that $R \rightarrow -1$. As $\alpha \rightarrow 1$, $E[R] \rightarrow 1 - (1 - \gamma)/(1 - \gamma\epsilon)$. The searcher is always inclined to take the correct branch and the overall benefit of the hierarchy depends on the rate ϵ at which number of documents reduces as the searcher descends the hierarchy, with a maximum reward of 1 in the limit as $\epsilon \rightarrow 1$.

Consider the special case where $\gamma = 1/3$ and $\alpha = 1/2$, then the searcher has no guidance on which path to choose at each level in the tree and $p_i = q_i = s_i = 1/3$. Now $E[R] = -1/(3 - \epsilon)$. A negative reward in $-\left[\frac{1}{2}, \frac{1}{3}\right]$ illustrates that unless the hierarchy provides some useful guidance to the user, it is better to exhaustively search all documents.

3 Choosing a navigation policy

The navigation policy depends on the information that is available to determine which path to take. Depending on the application, different types of node summaries might be available. If the documents are organised using topic modelling, for example, there may be a set of words that describe each node. Alternatively a set of prototype examples of the contents of each child might be presented.

For the evaluation measure, we assume that, given a target t , it is possible to compute the distance between t and a cluster of documents, C , and write this distance as $d(t, C)$. In many applications, a pairwise distance between documents is available: $d(t_1, t_2)$, $\forall t_1 t_2 \in D$, and the cluster distance can be based on these distances e.g.

$$d(t, C) = \frac{1}{|C|} \sum_{t' \in C} d(t, t').$$

Other possibilities include defining a distance based on the word distribution in each cluster (e.g., see [6, 7]), when topic modelling has been used to construct the hierarchy. In any case, we develop the navigation policy under the assumption the information available to the searcher is the set the values $d(t, C)$ available for each child node, C that can be reached from a parent node.

3.1 Policy: Greedy Descent to Threshold Depth

To develop a navigation policy, a criterion is required to decide whether to end the descent into the hierarchy at the current node. Furthermore, if a decision to explore deeper is made, a criterion to determine which child to explore is required. A practical real scenario is that the searcher has some search budget in mind, and is willing to examine a cluster of documents when the number of documents in the cluster is below a certain threshold. This results in the

following probability:

$$r_i(\pi_1) = \begin{cases} 1, & \text{if } n_i \leq thr \\ 0, & \text{otherwise} \end{cases}$$

Also, the choice of child depends deterministically on $d(t, C)$, with the searcher choosing the child that is closest to the target. Let $\{C_1, \dots, C_k\}$ be the available child nodes at a particular decision point, let $\mathbb{1}()$ correspond to the indicator function which is 1 when its boolean argument is true and zero otherwise and let $\ell = \arg \min_j d(t, C_j)$. Then

$$\begin{aligned} p_i(\pi_1) &= (1 - r_i(\pi))\mathbb{1}(t \in C_\ell) \\ q_i(\pi_1) &= (1 - r_i(\pi))\mathbb{1}(t \notin C_\ell). \end{aligned}$$

This policy represents a reasonable approximation to how searchers explore real-world hierarchies and therefore it is the primary policy that we propose to use in our measure. A crisp hierarchy in which the $d(t, C_j)$ are distinct and items tend to be organised in coherent clusters will evaluate highly in our measure with this policy. Outlying items that are equally distant from several clusters are likely to yield negative reward, as the searcher is likely to eventually choose an incorrect path, the deeper the descent into the hierarchy. The value of the threshold depth should be chosen based on the application context.

4 Toy Example and Conclusion

As an example, we generate $M = 10,000$ random points arranged in 8 clusters with samples spread around each cluster center as shown in Figure 1a. Figure 1b illustrates the binary tree that is obtained on application of agglomerative clustering using the pairwise distance between samples. Every node contains all items which appear in the lower levels, therefore we can decide in each step to: examine all the documents in the current node to retrieve our target, navigate to the child where the target is located, or choose the wrong path to others.

We apply our hierarchical evaluation following the policy described in 3.1 with a threshold of $n/8$. A score per item is obtained, measuring the quality of the hierarchy for this specific retrieved target.

As an overall evaluation we compute the mean value of these scores, obtaining a hierarchy quality $HQ = 0.52$. Examining the distribution of the quality reveals that 79% of the values are above 0.8 quality, according to the expected 1/8 reduction in search complexity when the searcher succeeds in finding the target. A minority of samples get a negative reward, as the hierarchy leads them to the incorrect node/child, reducing the overall quality to 0.52. Let's examine the targets with negative evaluation (see Figure 1a), these points belong to clusters 5 and 6. As Figure 1b shows, all points in cluster 5 and some in cluster 6 are closer on average to cluster 4 than they are to the full cluster formed by 1, 5, 6.




Fig. 1. Figure (a) shows the 10000 random samples generated in 8 clusters, and targets with negative evaluation in the black triangles. Figure (b) represents the hierarchy obtained with agglomerative clustering and the flaws in the dendrogram.

Therefore, when the navigation reaches that decision point in the hierarchy, it chooses the wrong path to cluster 4 and hence misses the target in cluster 5. So, although the hierarchy has chosen the right cluster in its leaf nodes, it contains a flaw at that internal node. Our proposed measure captures that flaw successfully.

Finally, we can conclude the proposed measure provides a reliable measure of the hierarchical quality and successfully detects flaws in the dendrogram. Further analysis should be done in a wide range of situations and different policies should be applied.

Acknowledgments

This project has been funded by Science Foundation Ireland under Grant No. SFI/12/RC/2289.

References

1. Y. C. Liu, Z. M. Li, H. Xiong, X. D. Gao, J. J. Wu, and S. Wu, "Understanding and enhancement of internal clustering validation measures," *IEEE Transactions on Cybernetics*, vol. 43, no. 3, pp. 982–994, 2013.
2. J. M. Cigarran, A. Penas, J. Gonzalo, and F. Verdejo, "Evaluating hierarchical clustering of search results," in *SPIRE'05*.
3. D. M. Johnson, C. Xiong, J. Gao, and J. J. Corso, "Comprehensive cross-hierarchy cluster agreement evaluation," in *AAAI (Late-Breaking Developments)*, 2013.
4. A. Gosavi, "Reinforcement learning: A tutorial survey and recent advances," *INFORMS Journal on Computing*, vol. 21, no. 2, pp. 178–192, 2009.
5. R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, vol. 1. MIT press Cambridge, 1998.
6. D. Kuang and H. Park, "Fast rank-2 nonnegative matrix factorization for hierarchical document clustering," in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 739–747, ACM, 2013.
7. D. M. Blei, M. I. Jordan, T. L. Griffiths, and J. B. Tenenbaum, "Hierarchical topic models and the nested chinese restaurant process," in *NIPS'03*.