Learning to Fingerprint the Latent Structure in Question Articulation

Kumar Mrityunjay[+] and Guntur Ravindra[++]
Technology Excellence Group[+], Senior Member of the ACM[++]
Email: mjay.cse@gmail.com[+], gravindra@acm.org[++]
Talentica Software, Pune, India

Abstract

Machine understanding of questions is tightly related to recognition of articulation in the context of the computational capabilities of an underlying processing algorithm. In this paper a mathematical model to capture and distinguish the latent structure in the articulation of questions is presented. We propose an objective-driven approach to represent this latent structure and show that such an approach is beneficial when examples of complementary objectives are not available. We show that the latent structure can be represented as a system that maximizes a cost function related to the underlying objective. Further, we show that the optimization formulation can be approximated to building a memory of patterns represented as a trained neural auto-encoder. Experimental evaluation using many clusters of questions, each related to an objective, shows 80% recognition accuracy and negligible false positive across these clusters of questions. We then extend the same memory to a related task where the goal is to iteratively refine a dataset of questions based on the latent articulation. We also demonstrate a refinement scheme called K-fingerprints, that achieves nearly 100% recognition with negligible false positive across the different clusters of questions.

1 Introduction

In practical applications, algorithms that power natural language interfaces designed to answer questions have to grapple with various types of user inputs and articulation styles. Inputs could be grammatically incorrect, they could be in a natural language that the algorithms are not trained for, and they may be too complex for the current algorithm to interpret. Despite challenging input variations, these interfaces are expected to respond back even if they are unable to interpret the input. Interpretation is always in the context of an underlying algorithm that performs the task of answer generation. The complexity of this underlying algorithm can vary depending on answer data as well as question interpretation capabilities in relation to the answer data. We refer to these aspects of answerability as an "objective". We argue that question answerability is always in the context of the objective. As a result, we need a system that can find a mapping between the latent structure of articulation and the algorithm that processes the question for a underlying knowledge-base.

1.1 Motivation

The challenging input conditions that question-answering systems have to grapple with, motivates us to explore how pattern recognition techniques may be used to detect if input questions have latent patterns that are friendly to an objective. The natural language interface then uses this pattern recognition system to detect if the input question matches the objective, and if it does, the objective is triggered. In practical applications we only have data that can be friendly to an objective, but the complementary dataset is a vast universe of possibilities for which we do not have complete information. Hence the system needs to account for the complementary information in the pattern recognition system without having explicit information of this complementary set. In this paper we address this issue by first formulating a mathematical model for the same, and then approximating the model's implementation to that of a neural auto-encoder. We then show how the same model can be used to iteratively learn and re-partition a data set of questions based on a latent structure in the articulation style.

The rest of this paper is organized as follows: in section-2 we present a literature review covering techniques for question analysis in some applications. In section-3 a theoretical formulation for the proposed pattern recognition system is presented. An iterative learning process called K-fingerprints is detailed in section-4 followed by experimental evaluation in section-5, and a conclusion in section-6.

2 Literature Review

Analysis of question structure and its impact on performance of answering systems has been well studied in the context of natural language knowledge repositories [1]. The general trend in question classification has been to use the nature of the answer classes , and identification of a factual intent (wh-type) [2]. A large class of question answering techniques tie the question domain to the knowledge domain and its representation. In recent years there has been growing interest in analyzing question structure in the context of answerability by humans. In [3] the authors present a supervised learning approach to predict if a question shall be answered or not. They propose use of a trained SVM classifier where multiple features such as length of the question, n-grams matching in a known corpus, POS diversity, LDA topic diversity, question editing activities, and characteristics of topic hierarchy are used to determine answerability. In another research [4], which also used a SVM classifier, a mix of boolean features, counting-based features and scoring-based features were used to classify questions in terms of answerability. The authors also analyzed answerability in terms of classifier performance on question types based on elicited response types such as advise, fact, opinion. In a related work [5], detection of questions that can go unanswered using heuristic features such as question length, asker history, polite words, and question subjectivity was presented. In [6], question answerability was modeled using a combination of title tokens, body tokens, sentiment, user data, and time. In summary, the goals of popular research in question answerability is to be able to model the relationship between question structure and its answerability so as to predict the answerability itself.

Contrary to popular literature, techniques presented in this paper have a slightly different goal, although related to answerability of questions. An underlying natural language processing algorithm designed to function deterministically for a specific objective works well when the latent structure

in question articulation matches what the algorithm has been designed for. Hence here we are interested in representing and detecting the latent structure in articulation of a question.

One possible approach to implement such a detector is to build a set of Bloomfilters [7] each representing one objective. Every question is now passed to all the Bloomfilters, and the filter that recognizes the pattern in the question identifies the question with the corresponding objective. However, one would have to still handle false positives as multiple Bloomfilters could mark the question as detected. Further, the choice of number of bits in the Bloomfilter depends on expected false positive which in-turn can de determined if all the possible variations in question articulation is available apriori. In cases where we have small application specific data this requirement cannot be satisfied. Yet another possibility is to use locality sensitive hashing [8, 9] where questions with similar latent patterns can be hashed to the same hash bucket. Detection would then entail hashing an input question's latent patterns to a bucket, and selecting the label assigned to the bucket as the objective. However, in order to achieve high recall and zero false positive one would have to use large number of hash functions and account for unseen data. Further, hashing based schemes fundamentally do not account for any non-linearities in the latent structure of articulation when such a structure can be numerically transformed.

3 Recognizing inputs based on an objective

In the context of question style recognition, we define an objective $\mathcal{O}$ as something that can successfully determine if the question that is being asked can be answered by the underlying system. Recognition of an objective depends on the ability to algorithmically understand the articulation of a question. The underlying recognition system also has data requirements. For example if the user asks a question related to Soccer, the system can generate answers only if data related to Soccer is present in the data store. As another example, if the data is stored in a relational database, the system can generate an answer via a database query only if it is capable of figuring out the correct database columns from the words in the question. Yet another example is a system that employs an algorithm that can convert questions starting with "how many" into a SQL query, but the same algorithm could generate the wrong query for any other type of questions. Hence defining the objective $\mathcal{O}$ is a key starting point.

Objectives that are algorithm and data dependent are very difficult to model as a computational entity. Hence we propose a workflow that involves (1) representing questions as feature vectors (2) clustering questions based on feature vectors (3) human evaluation of sampled questions from clusters (4) selection of question clusters that match $\mathcal{O}$ (5) modeling latent patterns in articulation of questions as a fingerprint (6) using fingerprints to detect if a question matches $\mathcal{O}$. In order to implement these steps we need a formal approach to defining the problem.

3.1 Problem formulation

Let $Q = \{q_1, q_2, ..., q_n\}$ be a set of questions. We define a transformation $\mathcal{T}(q_i) \rightarrow s_i$ where s_i is a mathematical representation of q_i. The set $S = \{s_1, s_2, ..., s_n\}$ is the corresponding mathematical representation of Q. We define an open set $S^{'} \subseteq S$ where members of $S^{'}$ meet the objective $\mathcal{O}$. We are interested in a function $f(*)$ such that

$$f(s_i) \rightarrow \mathbb{R}_i^d \forall s_i \in S'$$
$$f^{-1}(\mathbb{R}_i^d) \in S' \tag{3.1}$$

Such a function can transform any question into a d dimensional feature vector with a constraint that an inverse transformation would result in one of the elements in S'. An identity function could have satisfied such a relationship, but this is not very useful as questions unrelated to $\mathcal{O}$ would also be wrongly identified as being related to $\mathcal{O}$. Hence we need an additional constraint on $f(*)$ defined as

$$f^{-1}(\mathbb{R}_k^d) \notin S' \forall s_k \in S - S' \tag{3.2}$$

In (3.2), we say that any question from outside the set S' would result in an inverse that does not belong to S' and therefore does not match $\mathcal{O}$. In most practical applications where $\mathcal{O}$ is known, one can determine S' easily. But the complementary set $\widetilde{S'} = S - S'$ is unknown and difficult to find. Hence we need to find the $f(*)$ without the knowledge of $\widetilde{S'}$. However, before we can discover $f(*)$ we need to define $\mathcal{T}$.

3.2 Defining the transformation function $\mathcal{T}$

We need to define the transformation $\mathcal{T}(q_i)$ so that a question in natural language can be converted to a form suitable for computational manipulation. The algorithmic steps for this transformation are shown in algo-(1). We believe that this approach is most suitable when one wants to identify questions that are relevant for a given $\mathcal{O}$.

One can model a question as a sequence of words from a finite vocabulary. However, specific words and word orderings are used while expressing information as a question. When we generate a parse tree for a question ($PARSE(q_i)$), the tree is a representation of the underlying structure in expressing information as a question. However, there can be different trees for different questions, and questions of similar types have similar trees. The notion of similarity in questions is dependent on $\mathcal{O}$. Example of parse trees is shown in figure-3.1. In order to be able to compare the structure of trees, we use a string-matching approach. We traverse the nodes ($TRAVERSE(T_i)$) of the parse tree to generate a sequence of node labels. This sequence is represented as partial orders with a known rule ($SEQUENCES(B_i)$). For example, one could convert a sequence into non-overlapping segments of fixed length by skipping even indices. In addition one could create additional segments by skipping odd indices. For example, in figure-3.1 we have three questions that have different parse trees. Two of the questions have a similar parse tree. If we perform a breadth-first search and aggregate trigrams of parse nodes, we get sets of trigrams that are similar across all the trees. However that last tree has additional trigrams that are different from those for the first two trees. This difference can be captured as a difference in articulation.

Using this information a binary feature vector is created such that existence of a particular symbol/sequence is represented by 1 and 0 otherwise. Algo-1 describes the various steps involved in converting a set of questions into the mathematical representation referred to in (3.1). On line-11 in algo-1, the variable $\tau(x)$ is a weight. This weight can be an importance value derived relative to all the other $x \in D$. It could also be an indicator variable that takes a value 1. In this paper we use the indicator variable version.

How many students are registered for this course

['S_WRB_NN', 'WRB_NN_NP', 'NN_NP_VP', 'NP_VP_VP', 'VP_VP_NN', 'VP_NN_NP', 'NN_NP_NP', 'NP_NP_NP']

How many buses are plying

['S_WRB_NN', 'WRB_NN_NP', 'NN_NP_VP', 'NP_VP_VP', 'VP_VP_NP']

When Alexander reached India who stopped his further advances

['S_WRB_NN', 'WRB_NN_NP', 'NN_NP_VP', 'NP_VP_NN', 'VP_NN_NP', 'NN_NP_WP', **'NP_WP_VP'**, **'WP_VP_NN'**, 'VP_NN_NP', 'NN_NP_NP', 'NP_NP_NP']
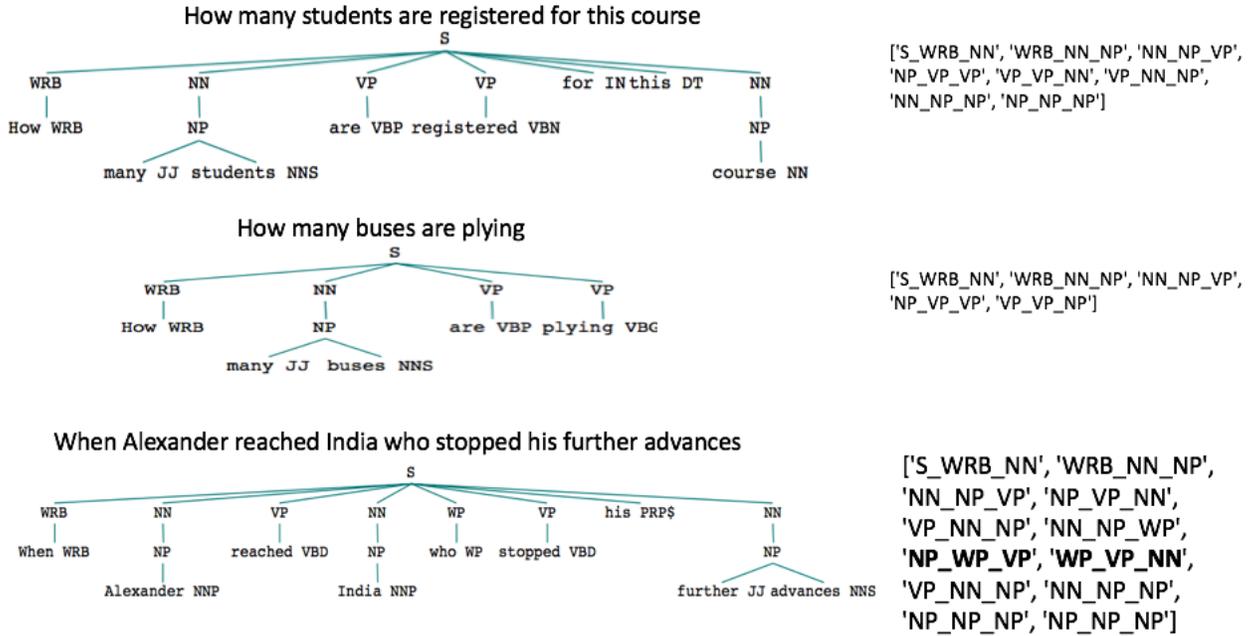
Figure 3.1: Parse trees for different questions and their traversal sequences

---

Algorithm 1 Transforming questions into vectors

Require: Q is a non-empty set
1: $D \leftarrow \{\}$
2: **for all** $q_i \in Q$ **do**
3:     $T_i \leftarrow PARSE(q_i)$
4:     $B_i \leftarrow TRAVERSE(T_i)$
5:     $C_i \leftarrow SEQUENCES(B_i)$
6:     $D \leftarrow D \cup C_i$
7: **end for**
Ensure: D is a non-empty and a unique set of elements
Ensure: D is an ordered set
Ensure: $d = |D|$
Ensure: $INDEX(x, D)$ is the position of x in the ordered set D
8: **for all** $q_i \in Q$ **do**
9:     $\mathbb{R}^d = 0$
10:     **for all** $x \in C_i$ **do**
11:         $\mathbb{R}[INDEX(x, D)] = \tau(x)$
12:     **end for**
13:     $\mathcal{T}(q_i) \rightarrow s_i = \mathbb{R}^d$
14: **end for**

Algorithm 2 Defining the similarity function $SIM()$

Ensure: $\tau = 1$
1: X is the input to the function f
2: $Y = f^{-1}(f(X))$
3: $ILARGEST([], n)$ takes vector as input and returns index of n largest values
4: $A \leftarrow$ indices i of X where $X(i) = \tau$
5: $B \leftarrow ILARGEST(Y, |A|)$
6: $SIM \leftarrow \frac{|A \cap B|}{|A \cup B|}$

3.3   Discovering $f(*)$

The asymmetric function $f(*)$ in (3.1) is a one-to-one transformation where as its inverse (3.2) is a one-to-many transformation. If we find a function whose inverse is also one-to-one transformation, then an identity function would have sufficed. But an identity cannot account for the additional constraint (3.2). Interestingly, we need to determine this function without the explicit knowledge of the complementary set $\widetilde{S}'$. In order to make the function discovery feasible we employ two computational strategies. We assume a universe of functions $\mathcal{F}$, of which $f^* \in \mathcal{F}$ is the optimal function satisfying (3.1) and the constraint (3.2). Secondly, to satisfy the constraint (3.2) without the knowledge of the complementary set $\widetilde{S}'$, we relax the one-to-one transformation to a many-to-one transformation of a reduced dimensionality. This means $f(s_i)$ results in a vector of dimensionality much lower than that of s_i. The inverse continues to be a one-to-many transformation where we constrain the transformation to any element in set S'. As a result of this relaxation we can then express the function discovery as a mathematical expression shown in (3.3)

$$\arg\max_{f \in \mathcal{F}} \frac{1}{|S'|} \sum_{s_i \in S'} \max_{s_k \in S'} \left(SIM \left[f^{-1}(f(s_i)), s_k \right] \right) \tag{3.3}$$

where, SIM is a measure of similarity.

3.4   Practical implementation as a fingerprinting system

The optimization function in (3.3) can be implemented as an encoder-decoder formulation where f represents the encoder and f^{-1} is the decoder. The encoder-decoder combine needs to be jointly optimized in order to implement (3.3). The transformation $\mathcal{T}$ described earlier converts each question into a vector of weights of a fixed dimensionality. Hence one of the frameworks that can jointly optimize the encoder-decoder combine is a neural auto-encoder with multiple hidden layers. We use an eight layer auto-encoder with Tanh as the activation function for each hidden layer, whereas one hidden layer has a ReLU activation [10]. The output layer has linear activation. The chosen Adam optimizer [11] runs for 300 epochs with a learning rate of 0.001 and mean square error as the loss function. We then define the similarity function SIM mentioned in (3.3) as shown in algo-2. This function measures the detection accuracy on a scale of 0 to 1. A perfect reconstruction of the input in the output layer of a trained auto-encoder would qualify for a detection accuracy of 100% and algo-2 guarantees this.

In the remainder of the paper we use the terms fingerprinting system and auto-encoder interchangeably as the fingerprinting technique is executed by a trained auto-encoder.

4 K-fingerprints

The fingerprinting system described in the previous section can be used to iteratively refine clusters by learning new latent patterns. Further, it can be used to refine clusters of similarly articulated questions by eliminating members that do not confirm to a pattern. We refer to this technique as K-fingerprints as we use a set of K trained auto-encoders $AE = \{ae_1^t, ae_2^t, ..., ae_K^t\}$. Steps implemented in this algorithm are summarized in algo-3.

To begin with, a large dataset of questions is first clustered into K clusters $\hat{C} = \{\hat{C}_1, \hat{C}_2, ..., \hat{C}_K\}$ using an algorithm like K-means. Owing to inaccuracy in the clustering algorithm, it is possible that some of the clusters do not have questions with a common latent pattern. At time step $t = 0$, a random sample of questions is selected from each cluster to create a base set of K clusters $C = \{C_1^0, C_2^0, ..., C_K^0\}$ whose latent patterns are to be learned by the K auto-encoders (one per cluster). We also maintain a set of cluster cardinalities $M = \{m_1^0, m_2^0, ..., m_K^0\}$ and these are updated in every iteration.

4.1   Iterative re-training with filtering

At $t = 0$, $C_i^{t=0}$ is used to train the auto-encoder $ae_i^{t=0}$ resulting in baseline auto-encoders. The baseline auto-encoders are iteratively retrained over multiple time steps until a convergence criterion is met. Details of iterative training is given in algo-3. The iterative training process involves two major phases, namely $DETECT$, and $TRAIN$. In the $DETECT$ phase at step t, questions from $\hat{C}_i$ are passed as input to ae_i^{t-1}. In time step $t-1$ the auto-encoder ae_i^{t-1} would have been trained using questions from C_i^{t-1} of cardinality m_i^{t-1}. At time-step t, each auto-encoder ae_i^{t-1} is made to detect all possible questions in $\hat{C}_i$. Say ae_i^{t-1} detects a set of $C_i^t \subseteq \hat{C}_i$ questions. If the cardinality of C_i^t is not the same as m_i^{t-1} then the i^{th} auto-encoder survives the current round. If there is no change, then the i^{th} auto-encoder is eliminated from participation in subsequent rounds and is added to another set F. All the auto-encoders that have survived the next round, now move on to the $TRAIN$ phase. In the $TRAIN$ phase, auto-encoder ae_i^{t-1} is trained using questions from C_i^t to result in ae_i^t and M is updated to $M = \{m_1^t, m_2^t, ..., m_K^t\}$ where $m_i^t = |C_i^t|$. This process is repeated until F ends up with K trained auto-encoders or we have looped through the process enough number of times. After the iteration loop ends, auto-encoders that are not part of F now become part of F provided $m_i^t > 0$.

5   Experiments and Results

Approaches described in the previous sections are evaluated under two broad criteria. On one hand we want the fingerprint system to detect questions with articulation that matches its memory. At the same time, we want each fingerprint system to fail on detecting questions used to train the other fingerprint systems. The second criterion is harder to meet as there are $K-1$ complementary memories for every fingerprint system and there can be some questions that have an articulation which could match with multiple memories. In other words there can be confusion among multiple fingerprint systems. In applications such as chat-bots or translation systems we require the fingerprint system to have 0% false positives and a very high true positive rate in detecting the articulation type. In these applications a false positive would result in wrong interpretation or an

Algorithm 3 Training K-fingerprints

Require: $\hat{C} = \{\hat{C}_1, \hat{C}_2, ..., \hat{C}_K\}$ a set of K clusters of questions
Require: $C = \{C_1^0, C_2^0, ..., C_K^0\}$ a set of K clusters where $C_i^0 \subset \hat{C}_i$
Require: $AE = \{ae_1^0, ae_2^0, ..., ae_K^0\}$ a set of K baseline auto-encoders trained with C
Require: $M = \{m_1^0, m_2^0, ..., m_K^0\}$ a set of integers showing how many questions were detected by the baseline auto-encoders
Require: F is the final set of trained auto-encoders

 1: $G \leftarrow AE$
 2: Initialize iteration round $t \leftarrow 1$
 3: **while** G is not empty **do**
 4:      $G^t = \{\phi\}$, $C^t = \{\phi\}$, $M^t = \{\phi\}$
 5:      **for** Each auto-encoder AE_i^{t-1} in G **do**
 6:          $C_i^t \leftarrow \textbf{DETECT}(\textbf{AE}_\textbf{i}^{\textbf{t-1}}, \hat{\textbf{C}}_\textbf{i})$
 7:          **Continue** if C_i^t is empty
 8:          **if** items detected $|C_i^t| \neq m_i^{t-1}$ **then**
 9:              Include AE_i^{t-1} in G^t
10:              Include C_i^t in C^t
11:              Include $|C_i^t|$ in M^t
12:          **else**
13:              Include AE_i^{t-1} into F
14:          **end if**
15:      **end for**
16:      $M \leftarrow M^t$, $C \leftarrow C^t$, $AE = \{\phi\}$
17:      **for** $AE_i \in G^t$ **do**
18:          $AE_i^t \leftarrow \textbf{TRAIN}(\textbf{AE}_\textbf{i}, \textbf{C}_\textbf{i}^\textbf{t})$
19:          Include updated AE_i^t into AE
20:      **end for**
21:      $G \leftarrow AE$ (Updated list of auto-encoders)
22:      **Loop counter to exit**
23:      $t \leftarrow t + 1$ (Next Iteration)
24: **end while**
25: $F \leftarrow F \cup C$ (Final list of trained auto-encoders)
26: Use each AE_i in F and detect classes of questions in $\hat{C}_i$ using the $DETECT()$ function

Cluster	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14
Number of Questions	605	539	432	395	1269	1001	1304	889	509	531	1023	497	788	776

Figure 5.1: Number of questions in each cluster

out-of-context answer where as a low true positive would result in frequent no-response situations. Unlike typical machine learning applications, here we are not interested in unseen data (test data). We are interested in seeing if a pattern match system learned the known set of patterns well. Keeping this type of scenario in mind we evaluate the fingerprinting system in terms of true positive rate and false positive rate.

5.1 Data set

To evaluate the approaches described in the previous sections, SQuAD data set [12] (dev-data set version-1.1) was used as a source of questions with various types of articulations. There were around 10k questions spanning 48 different topics. Questions were transformed into feature vectors using the transformation process described in section-3.2. For the tree-traversal, we used a breadth-first-search over the parse tree and converted sequence of parse labels into trigrams and 4-grams. We used an overlap of 2 and 3 for trigrams and 4-grams respectively.

This resulted in a dimensionality of 192 unique symbols (algo-1). Each question in turn, was transformed into a 192 dimensional feature vector where each appearing symbol was given a weight of $\tau = 1$. The questions in the feature vector form were clustered into 14 clusters using K-means++ [13]. The choice of 14 was based on inspecting the density of points in bins of a K-d tree. The goal of clustering is to help a human evaluator to identify consistent sets of clusters by randomly sampling and viewing questions in each cluster. The expectation is that questions in each cluster correspond to a specific type of articulation, and therefore different machine interpretability. Figure-5.1 shows the number of questions in each of the 14 clusters. We observe that some clusters have many more questions than others but there are at least 350 questions per cluster.

5.2 Baseline evaluation

From each of the 14 clusters 40 questions were selected per cluster at random. These questions were used to train 14 fingerprint systems (one per cluster) to get the baseline fingerprint systems. Note that the dimension of the input vector varies for each fingerprint system, and it depends on the total number of unique symbols in each cluster. Each baseline fingerprint system was then made to detect all the questions in the corresponding cluster. The goal was to observe how well the fingerprint system trained with a small set of questions can detect all the remaining questions in a cluster. Table-1 shows the number of questions detected by each baseline fingerprint system for the corresponding cluster. We observe that many fingerprint systems have a poor recall. This indicates that the corresponding clusters did not have questions with similar articulation. If the articulation would have been the same, we expect the feature vectors to be the same, and therefore a higher recall. The issue could be with poor clustering, and is resolved using K-fingerprints as explained later.

After we exclude clusters with low recall we have 6 clusters namely $\{C_1, C_2, C_4, C_6, C_{12}, C_{14}\}$. The questions in each of these selected clusters were used to train a corresponding fingerprint

Cluster	Questions	Trained with	Detected	Recall
C1	605	40	428	0.7074
C2	539	40	207	0.384
C3	432	40	0	0
C4	395	40	237	0.6
C5	1269	40	0	0
C6	1001	40	953	0.952
C7	1304	40	1	0.0008
C8	889	40	0	0
C9	509	40	34	0.0668
C10	531	40	106	0.1996
C11	1023	40	253	0.2473
C12	497	40	393	0.7907
C13	788	40	0	0
C14	776	40	572	0.7371

Table 1: Accuracy of baseline fingerprint systems

(a) $SIM>=0.5$ (b) $SIM>=0.7$ (c) $SIM>=0.9$
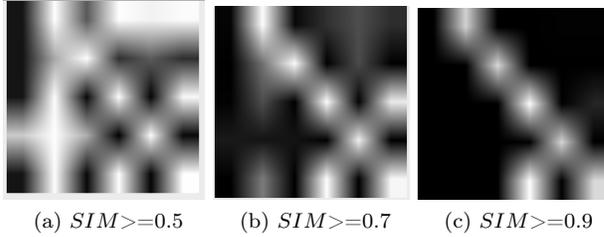
Figure 5.2: Confusion matrix for different thresholds viewed as a heat-map

system with parameters as described in section-3.4. Note that each fingerprint system learns to detect questions corresponding to the specific objective as represented by the articulations in the corresponding cluster. The trained fingerprint systems are now used to detect questions from the 6 clusters. We are interested in true positive for fingerprint system ae_i when questions in C_i are used as an input and the false positive when questions from other than C_i are used as an input to ae_i.

5.3   Evaluation

In order to evaluate the accuracy of the fingerprinting system, a 6x6 confusion matrix was generated corresponding to the selected clusters. Each (i, j) cell in the confusion matrix refers to the percentage of questions detected as relevant to the i^{th} cluster by the i^{th} fingerprinting system, when questions from the j^{th} objective/cluster were given as an input. Ideally one would like this value to be 0 when $i \neq j$ and 1 when $i = j$. Figure-5.2 shows the confusion matrix viewed as a heat-map when different thresholds were used on the $SIM()$ function described in algo-2. Note how the accuracy improves as the threshold increases from 0.5 to 0.9. Ideally we would like to get a bright line along the principal diagonal when the confusion is zero. We choose the threshold as 0.9 as it imposes hard performance limits in terms of confusion as well as true positives. Usually when the threshold is high the confusion reduces but the true positive rate also reduces. In the present case

	C1	C2	C4	C6	C12	C14
C1	0.793	0	0	0	0	0
C2	0	0.82	0.003	0	0.01	0.01
C4	0	0	0.853	0	0	0
C6	0	0	0	0.998	0	0.151
C12	0	0	0	0	0.845	0
C14	0	0	0	0.998	0	0.852

(a) Recognition accuracy and confusion matrix for selected clusters
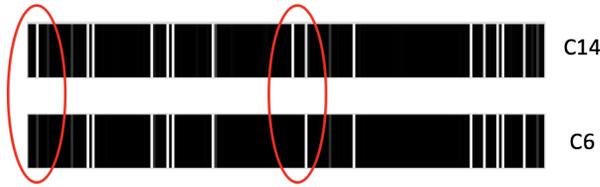
(b) Fingerprints showing that $C6$ and $C14$ are very similar with $C14$ being a superset of $C6$

Figure 5.3: Confusion matrix for the 6 clusters and cluster similarity for $C14$ and $C6$

			CLUSTERS																										
		C1		C2		C3		C4		C5		C6		C7		C8		C9		C10		C11		C12		C13		C14	
	C1	1	508/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0.1	34/357	0	0/415	0	0/0	0	0/581
	C2	0	0/508	1	324/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	1/415	0	0/0	0	0/581
C	C3	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
L	C4	0	0/508	0	0/324	0	0/0	1	242/242	0	0/746	0	0/1000	0.21	43173	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
U	C5	0	0/508	0	0/324	0	0/0	0	0/242	1	746/746	0	29/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	15/581
S	C6	0	0/508	0	0/324	0	0/0	0	0/242	0	19/746	1	1000/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	10/581
T	C7	0	0/508	0	0/324	0	0/0	0	3/242	0	0/746	0	0/1000	1	14/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
E	C8	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
R	C9	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	1	40/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
S	C10	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	1	106/106	0	0/357	0	0/415	0	0/0	0	0/581
	C11	0.1	38/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	1	357/357	0	0/415	0	0/0	0	0/581
	C12	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	1	415/415	0	0/0	0	0/581
	C13	0	0/508	0	0/324	0	0/0	0	0/242	0	0/746	0	0/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	0	0/581
	C14	0	0/508	0	0/324	0	0/0	0	0/242	0	22/746	1	999/1000	0	0/14	0	0/0	0	0/40	0	0/106	0	0/357	0	0/415	0	0/0	1	581/581

Table 2: Recognition accuracy, confusion matrix and cluster membership

we see good true positive rate as well as very low confusion (false positive across clusters).

Figure–5.3a shows the confusion matrix when the threshold on $SIM()$ is 0.9. We observe that all the fingerprints performed well in terms of confusion, except the one corresponding to $C14$. Here, $C14$ was confused with $C6$. In order to understand why $C14$ was confused with $C6$ but not vice versa, we visually plot the average of the input feature vector as an image as shown in 5.3b. We observe that the feature vector is very similar between $C6$ and $C14$, with $C14$ having two additional symbols ($C14$ is a superset of $C6$). This in turn implies that the clustering algorithm failed to group questions corresponding to $C6$ and $C14$ into a single cluster. However, the fingerprinting system detected the fact that $C6$ is similar to $C14$ and not vice versa there by justifying the clustering algorithm's behavior.

The performance of fingerprinting in terms of confusion matrix is very good, and the recall of most fingerprint systems is around 80%. A system of this form is very useful in use-cases where false positives cannot be tolerated where as false negatives can still be tolerated. Some typical use-cases are like chat-bots, where it is fine if the bot refuses to answer the question due to the verdict of the fingerprint system (false negative), but it is hilarious when an unrelated answer is generated in response to the question (false positive). The fingerprint system can avoid such a scenario.

5.4   K-fingerprints

As detailed in the previous sections, K-fingerprints is an approach to iteratively train and partition data so as to improve fingerprinting accuracy. In table-1 we observed that many clusters had a very poor recall, main reason being poor clustering accuracy. We use K-fingerprints to learn to refine clusters after the K auto-encoders are initially trained with a randomly selected set of 40 questions per cluster. During the iterative process each ae_i^{t-1} eliminates questions from the i^{th} cluster if the fingerprint for these questions do not match with the memory of ae_i^{t-1}. By refining the set iteratively, ae_i^{t-1} also learns new symbols and this appears as a change in input layer size, where retraining would result in a new memory for ae_i^t. When questions in a cluster are eliminated in a particular round, they may be included in another round of training and the repetitive $TRAIN$ and $DETECT$ steps described in algo-3 ensures this process.

In the present case we have 14 auto-encoders that iteratively refine the initial 14 clusters of questions. Table-2 summarizes multiple observations after K-fingerprints have been trained. The table shows the true positive recognition rate highlighted in green color. We observe that the true positive is nearly 100% in all the 14 cases and this is a marked improvement from the previous case where we dropped a few clusters due to clustering inconsistency. What K-fingerprints did was to eliminate some questions from each cluster so that it learns a consistent pattern from among the questions that remained in the cluster. We can observe that the number of questions in some of the clusters has actually reduced. For example $C1$ and $C11$ initially had 605 and 1023 questions respectively (table-1). However, in table-2 we see that the number of questions has reduced to 508 and 357 respectively. In addition, the detection accuracy is 100% for $C1$ and $C11$ where as cluster $C11$ was earlier dropped out of contention in the initial experiment due to poor clustering accuracy. In case of $C11$, K-fingerprints eliminated a large set of questions in order to create a consistent cluster for which the true positive rate is now 100%.

Other interesting case is with respect to $C7$ and $C9$. Here K-fingerprints reduced the cluster sizes by a significant number in order to improve the true positive (better cluster consistency)

without impacting the false positives. Clusters $C3$ and $C8$ could not be improved from the previous experiment where they were eliminated. They continue to stay eliminated after K-fingerprints. However, $C5$ is interesting as there was 0% match in the previous experiment where as after K-fingerprints, the cluster size was reduced from 1269 questions to 746 questions resulting in a 100% recognition accuracy (true positive) and 0% confusion with questions from other clusters (false positive).

6   Conclusion and Future Work

In this paper the notion of objective-driven question style detection was presented as a mathematical model. We argued that style of articulation of questions can be captured in a pattern recognition system which we refer to as fingerprinting. We showed how the fingerprinting system can be implemented using auto-encoders, and explained why an auto-encoder is a suitable architecture to represent the mathematical model. Algorithms to convert groups of questions with similar articulation into a mathematical model of representation were described. Applications in the domains of chat bots and translation systems were cited and performance evaluation conducted using a dataset of over 10K questions.

The idea of fingerprinting was extended to what we referred to as K-fingerprints, where we showed how an iterative process can be used to refine clusters of questions in order to generate a fingerprint that can capture consistency among cluster members. We showed that the recognition accuracy is nearly 100% for each fingerprint with almost 0% error in recognizing fingerprints of unrelated clusters.

As an extension to the idea of K-fingerprints, one can now implement an iterative clustering scheme. Instead of eliminating questions that do not have a common articulation, we could extend K-fingerprints to move questions among clusters thereby creating newer clusters.

One can observe that once the questions were transformed using the transformation function, rest of the steps are outside the realm of natural language and text processing. This motivates us to explore the possibility of applying the algorithms described to other types of data such as images. We believe that the mathematics would remain the same but the computational method for function discovery would change. It would have to account for the spatial nature of information in the images. Adding convolutional layers in combination with dense layers in the auto-encoder would be beneficial. One needs to explore though.

References

[1] D.Moldovan, S.Harabagiu, and M.Pasca, "The structure and performance of an open-domain question answering system," in *Proceedings of the 38th Annual Meeting on Association for Computational Linguistics*, 2000, pp. 563–570.

[2] B.Fan, Z.Xingwei, and Y.Hau, "Function based question classification for general qa," in *Proceedings of Conference on Empirical Methods in Natural Language Processing*, 2010, pp. 1119–1128.

[3] S.K.Maity, A.Kharb, and A.Mukherjee, "Language use matters: Analysis of the linguistic structure of question texts can characterize answerability in quora," *CoRR*, 2017.

[4] C. Shah, V. Kitzie, and E. Choi, "Questioning the question – addressing the answerability of questions in community question-answering," in *47th Hawaii International Conference on System Sciences.* IEEE, Jan 2014, pp. 1386–1395.

[5] L.Yang, S.Bao, Q.Lin, X.Wu, D.Han, Z.Su, and Y.Yu, "Analyzing and predicting not-answered questions in community-based question answering services," in *AAAI'11*, 2011, pp. 1273–1278.

[6] G.Dror, Y.Maarek, and I.Szpektor, "Will my question be answered? predicting question answerability in community question-answering sites," in *ECML- PKDD '13*, 2013.

[7] B. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.

[8] P.Indyk and R.Motwani, "Approximate nearest neighbors: Towards removing the curse of dimensionality," in *Proceedings of the 30th ACM Symposium on Theory of Computing.* ACM, 1998, pp. 604–613.

[9] R.Panigrahy, "Entropy based nearest neighbor search in high dimensions," in *Proceedings of the 17th annual ACM-SIAM Symposium on Discrete Algorithms.* Society for Industrial and Applied Mathematics, 2006, pp. 1186–1195.

[10] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010, pp. 807–814.

[11] D. Kingma and J.Ba, "Adam: A method for stochastic optimization," in *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015, pp. 1–13.

[12] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang, "Squad: 100,000+ questions for machine comprehension of text," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 2383–2392.

[13] D.Arthur and S.Vassilvitskii, "K-means++: The advantages of careful seeding," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '07. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2007, pp. 1027–1035. [Online]. Available: http://dl.acm.org/citation.cfm?id=1283383.1283494