

Deep Learning on Attributed Sequences

by

Zhongfang Zhuang

A Dissertation

Submitted to the Faculty

of the

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Doctor of Philosophy

in

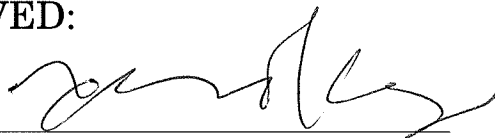
Computer Science

February 28, 2019

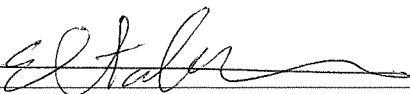
APPROVED:



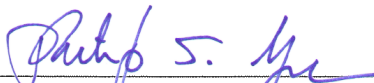
Professor Elke A. Rundensteiner
Worcester Polytechnic Institute
Advisor



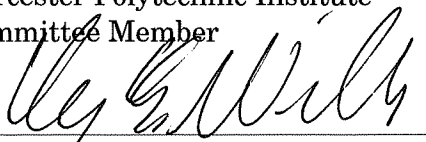
Professor Xiangnan Kong
Worcester Polytechnic Institute
Co-Advisor



Professor Mohamed Y. Eltabakh
Worcester Polytechnic Institute
Committee Member



Professor Philip S. Yu
University of Illinois at Chicago
External Committee Member



Professor Craig Wills
Worcester Polytechnic Institute
Head of Department

Contents

1 Introduction	13
1.1 The Prevalence of Attributed Sequences	14
1.2 Applications using Attributed Sequences	16
1.3 State-of-the-Art	21
1.3.1 Attributed Sequence Embedding	21
1.3.2 Deep Metric Learning on Attributed Sequences . .	22
1.3.3 One-shot Learning on Attributed Sequences	24
1.3.4 Attention for Attributed Sequence Classification . .	24
1.4 Research Challenges	25
1.4.1 Heterogeneous dependencies in attributed sequences	25
1.4.2 Lack of labeled data	26
1.4.3 Metric Learning on Attributed Sequences	27
1.4.4 Generalize from Only One Sample	27

1.4.5 Attention in Attributes and Sequences	28
1.5 Proposed Solutions	28
1.5.1 Attributed Sequence Embedding	29
1.5.2 Incorporating Feedback of Attributed Sequences . .	30
1.5.3 Classify Attributed Sequences in One-shot	32
1.5.4 Attention Model for Attributed Sequence Classifi- cation	33
1.6 Road Map	34
2 Background: Attributed Sequence	35
3 Unsupervised Attributed Sequence Embedding	39
3.1 Problem Definition	40
3.2 The NAS Framework	41
3.2.1 Attribute Network	42
3.2.2 Sequence Network	44
3.2.3 Attributed Sequence Learning	45
3.3 Experimental Setup	49
3.3.1 Data Collection	49
3.3.2 Compared Methods	50
3.3.3 Network Parameters	52

3.4	Outlier Detection Tasks	52
3.4.1	Performance	54
3.4.2	Parameter Study	56
3.5	Clustering Tasks	57
3.5.1	Performance	57
3.5.2	Parameter Study	59
3.5.3	Case Study: Security Management System	61
3.6	Scalability	64
4	Deep Metric Learning on Attributed Sequences	66
4.1	Problem Definition	67
4.2	The MLAS Network	68
4.2.1	AttNet and SeqNet	69
4.2.2	FusionNet	71
4.2.3	MetricNet	74
4.3	Learning Feedback	75
4.4	Experiments	78
4.4.1	Datasets	78
4.4.2	Compared Methods	80
4.4.3	Experimental Settings	82
4.4.4	Results and Analysis	85

4.4.5 Case Studies	92
5 One-shot Learning on Attributed Sequences	93
5.1 Problem Definition	94
5.2 The OLAS Model	95
5.2.1 Approach	95
5.2.2 OLAS Model Design	99
5.2.3 OLAS Model Training	105
5.2.4 Labeling Attributed Sequences	109
5.3 Experiments	110
5.3.1 Datasets	110
5.3.2 Compared Methods	112
5.3.3 Experiment Settings	114
6 Attention Model for Attributed Sequence Classification	120
6.1 Problem Definition.	121
6.2 Attributed Sequence Attention Mechanism	121
6.2.1 Network Components.	122
6.2.2 Attributed Sequence Classification	128
6.2.3 Training	129
6.3 Experiments	130

6.3.1 Datasets	130
6.3.2 Compared Methods	134
6.3.3 Experimental Setting	135
6.3.4 Accuracy Results	135
6.3.5 Parameter Sensitivity Analysis	136
7 Related Work	140
7.1 Deep Learning	141
7.2 One-shot Learning	142
7.3 Attention Network	143
7.4 Sequence Mining	143
7.5 Clickstream Analysis	144
7.6 Metric Learning	145
8 Conclusion and Future Work	147
8.1 Conclusion	147
8.2 Future Work	149

List of Figures

1.1 An Example of Attributed Sequences.	15
1.2 Fraud Detection using Attributed Sequence Embedding	16
1.3 Motivating example: attack detection in one-shot	18
1.4 Comparison of different embedding problems.	20
1.5 Distance metric learning on attributed sequences.	22
1.6 A comparison of related classification problems.	23
1.7 Three Types of Dependencies in An Attributed Sequence.	25
3.1 Process Flow of NAS	41
3.2 NAS performance of k-NN outlier detection	53
3.3 NAS parameter sensitivity to k values	54
3.4 NAS performance in outlier detection (varying epoches)	55
3.5 NAS performance in clustering	58
3.6 NAS parameter study in clustering	59

3.7 NAS parameter study with varying epochs	60
3.8 Case studies: NAS in clustering	61
3.9 NAS runtime scalability	64
4.1 Comparison of three different network designs.	72
4.2 MLAS Network with Balanced Design	76
4.3 The effect of feedback in MLAS	85
4.4 MLAS performance with varying clustering parameters	87
4.5 MLAS parameter study with different network dimensions	88
4.6 Effect of pre-training parameters in MLAS	89
4.7 Plots of MLAS feature space	90
5.1 OLAS network structure	99
5.2 OLAS accuracy on AMS data with the Euclidean distance	114
5.3 OLAS accuracy on Wikispeedia data with the Euclidean distance	115
5.4 OLAS accuracy on AMS data with the Manhattan distance	117
5.5 OLAS accuracy on Wikispeedia data with the Manhattan distance	118
6.1 Two types of attention for attributed sequence classification	123
6.2 AMAS performance evaluation	132

6.3 AMAS performance with adaptive training	133
6.4 Comparison of the history of training and validation losses.	137
6.5 Plot of weights in AMAS models	138

Abstract

Recent research in feature learning has been extended to sequence data, where each instance consists of a sequence of heterogeneous items with a variable length. However, in many real-world applications, the data exists in the form of attributed sequences, which is composed of a set of fixed-size attributes and variable-length sequences with dependencies between them. In the attributed sequence context, feature learning remains challenging due to the dependencies between sequences and their associated attributes. In this dissertation, we focus on analyzing and building deep learning models for four new problems on attributed sequences.

First, we propose a framework, called NAS, to produce feature representations of attributed sequences in an *unsupervised* fashion. The NAS is capable of producing task independent embeddings that can be used in various mining tasks of attributed sequences.

Second, we study the problem of deep metric learning on attributed sequences. The goal is to learn a distance metric based on *pairwise user feedback*. In this task, we propose a framework, called MLAS, to learn a distance metric that measures the similarity and dissimilarity between attributed sequence feedback pairs.

Third, we study the problem of one-shot learning on attributed sequences. This problem is important for a variety of real-world applications ranging from fraud prevention to network intrusion detection. We design a deep learning framework OLAS to tackle this problem. Once the OLAS is trained, we can then use it to make predictions for not only the new data but also for entire previously unseen new classes.

Lastly, we investigate the problem of attributed sequence classification with attention model. This is challenging that now we need to assess the importance of each item in each sequence considering both the sequence itself and the associated attributes. In this work, we propose a framework, called AMAS, to classify attributed sequences using the information from the sequences, metadata, and the computed attention.

Our extensive experiments on real-world datasets demonstrate that the proposed solutions significantly improve the performance of each task over the state-of-the-art methods on attributed sequences.

Acknowledgments

I would like to express my gratitude to my advisor Dr. Elke Rundensteiner and co-advisor Dr. Xiangnan Kong for their excellent guidance, patience, and support for my research. I would like to thank Dr. Mohamed Eltabakh for collaborating with me on the recurring query processing projects during my Ph.D. qualifier. My special thank you goes to Dr. Philip Yu for devoting his time and efforts to serve on my Ph.D. committee.

My sincere thank you also goes to the Computer Science Department at Worcester Polytechnic Institute for the generous Teaching Assistantship offers in the early years of my Ph.D. study. This Teaching Assistantship enables and encourages me to follow my passion for computer science. I am also grateful for the generosity of Amadeus IT Group for supporting my research through the Research Assistantship as well as access to real-world data and problems.

I sincerely thank my friend and collaborator, Dr. Chuan Lei, for sharing his experiences in research during our collaborations of the big data system research. My sincere thank you also goes to Dr. Jihane Zouaoui and Rémi Domingues from Amadeus IT Group for the collaborations in deep learning research.

My thank you also goes to current and alumni DSRG students Dr. Dongqing Xiao, Xiao Qin, Xinyue Liu, Mingrui Wei, and Dr. Chiyang Wang for the support and advice during my Ph.D. years.

Last but not least, I want to thank my family for unconditionally supporting me to pursue my Ph.D.

Chapter 1

Introduction

1.1 The Prevalence of Attributed Sequences

Sequential data arises naturally in a wide range of applications [4, 44, 62, 66]. Examples of sequential data include clickstreams of web users, purchase histories of online customers, and DNA sequences of genes. Different from conventional multidimensional data [51] and time series data [30], sequential data [75] are not represented as feature representations of continuous values, but as sequences of categorical items with variable-lengths.

The sequential data in many real-world applications also often comes with a set of data attributes depicting the context. In this work, we will call this type of data, where each instance has both sequential data and the attributes, *attributed sequences*.

For example, in online ticketing systems as shown in Figure 1.1, each user transaction includes both a sequence of user actions (*e.g.*, “login”, “search” and “pick seats”) and a set of attributes (*e.g.*, “user name”, “browser” and “IP address”) indicating the context of the transaction. In gene function analysis, each gene can be represented by both a DNA sequence and a set of attributes indicating the expression levels of the gene in different types of cells. In the area of web search, an attributed sequence is composed of a static user profile (*e.g.*, “geolocation”) and

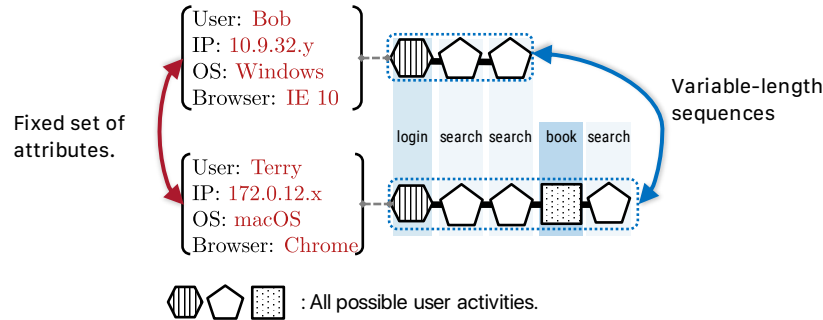


Figure 1.1: An example of attributed sequences. Each attributed sequence represents a transaction in an online ticketing system, which includes a *sequence* of user actions (e.g., “login”, “search”, “pick seats” and “confirm”) and a set of *attributes* for the transactional context.

the dynamic sequence of keywords searched (e.g., “snow storm” and “temperature”).

Many real-world applications involve mining tasks over the sequential data. For example, in online ticketing systems, administrators are interested in finding fraudulent sequences from the clickstreams of users [58, 66]. In user profiling systems, researchers are interested in grouping purchase histories of customers into clusters [62]. Motivated by these real-world applications, sequential data mining has received considerable attention in recent years [44, 4]. However, the attributes associated with these sequential data in these real-world applications have been overlooked despite their importance. Below, we highlight the importance of using attributed sequences in two real-world applications.

1.2 Applications using Attributed Sequences

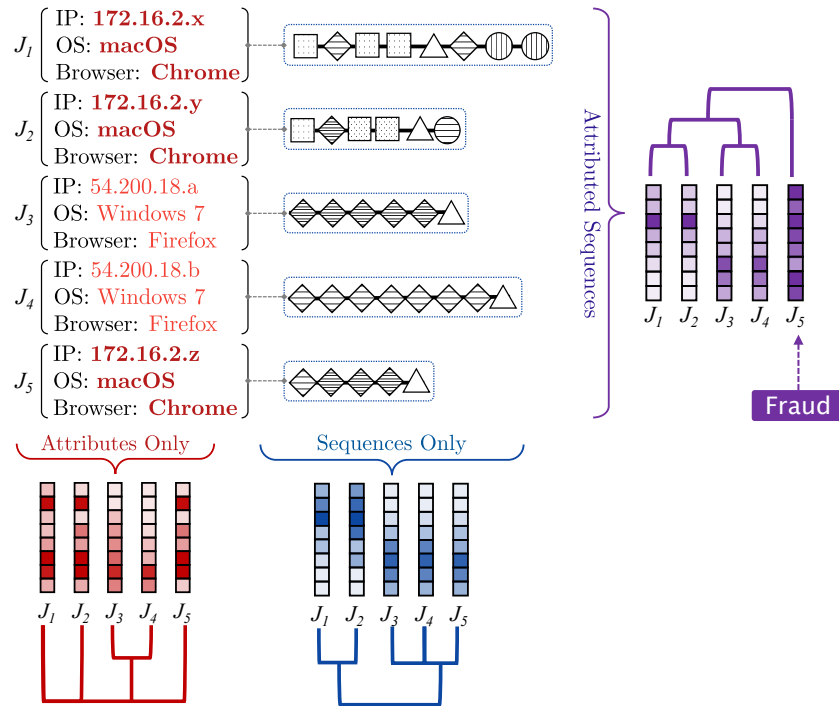


Figure 1.2: Dendrograms of feature representations learned from attributed sequences for fraud detection tasks. J_5 is a user committing fraud. However, it is considered a normal user session by the feature representation generated using either only attributes or only sequences. J_5 can only be caught as a fraud instance using the feature representation generated with regards to both attributes and sequences.

- Airline Ticket Booking Fraud.** In an airline ticketing system, whether a sequence of user actions is fraudulent depends closely upon the context of the transaction. A sequence of user actions, like “*apply business travel discount, then book a ticket*”, will be normal

for a *business travel* context, but could be fraudulent for a *leisure travel* context. In Figure 1.2, we present an example of fraud detection from a user privilege management system in Amadeus [2]. This system logs each user session as an attributed sequence (denoted as $J_1 \sim J_5$). Each attributed sequence consists of a sequence of user’s activities and a set of attributes derived from metadata values. The attributes (e.g., “IP”, “OS” and “Browser”) are recorded when a user logs into the system and remain unchanged during each user session. We use different shapes and colors to denote different user activities, e.g., “Reset password”, “Delete a user”. An important step in this fraud detection system is to “red flag” suspicious user sessions for potential security breaches. In Figure 1.2, we observe three groups of feature representations learned from the Amadeus application logs. For each group, we use a dendrogram to demonstrate the similarities between feature representations within that group. Neither of the feature representations using only sequences or only attributes detects any outliers due to not considering attribute-sequence dependencies. However, user session J_5 is discovered to be fraudulent using a learning algorithm that incorporates all three types of dependencies.

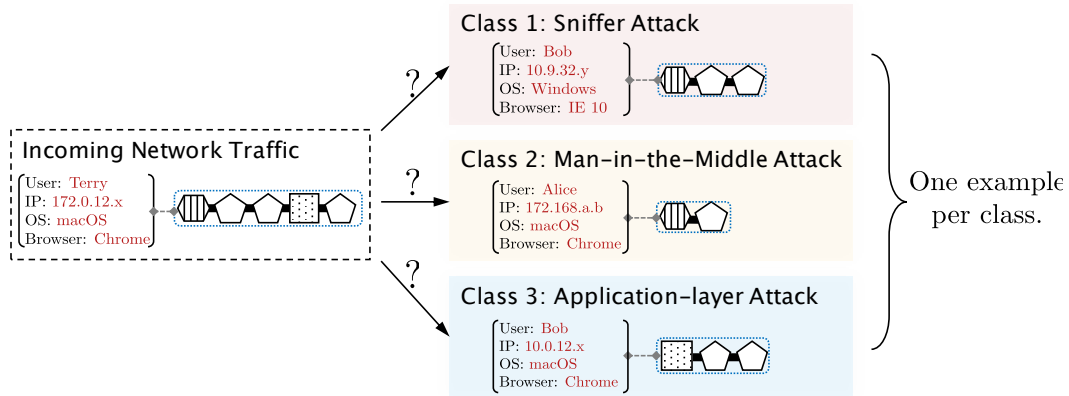


Figure 1.3: Network attack detection using one-shot learning on attributed sequences. Each instance is composed of a user profile as the *attributes* and a *sequence* of user actions (depicted using different shapes). A system administrator is interested in finding out if the incoming network traffic is malicious with **only one** sample per class.

2. **Bot Traffic Detection.** With the rapid advance in e-commerce, more businesses than ever in history are using websites to advertise their products and services. Attributed sequences exist naturally on these websites: the activities of each visitor are recorded in log files as *sequences* alongside the visitor's profile as *attributes*. Aside from the *real traffic* from potential customers on the website, there is another type of traffic, namely, the *bot traffic*. Bot traffic [69] is from a series of applications that run scripts over the websites for various task. Many of those tasks are malicious to the

websites, such as account hijacking with brute force, stealing web contents, undercutting prices and probing for potential attack opportunities [78]. Since the differences in attributes and sequences between bot and real traffic may not be significant, it is difficult to distill bot traffic from real traffic using only either attribute data or only sequence data. For instance, a real customer may have a similar or even identical profile (*e.g.*, “OS”, “IP”) as the bot scripts; the tasks (*e.g.*, search information, click menus) conducted by real customers and bot scripts may also be similar.

Thus, a bot traffic system should be able to (1) Utilize attributed sequences in the log files to infer the different patterns of real traffic and bot traffic and (2) Generalize the patterns of both types of traffic to prevent future bot traffic.

- 3. Network Intrusion Detection.** Network traffic can be modeled as attributed sequences. Namely, it consists of a sequence of packages being sent or received by the routers and a set of attributes indicating the context of the network traffic (*e.g.*, user privileges, security settings, *etc*). To respond in a timely fashion to potential network intrusion threats, one first has to determine what the intrusion type of incoming potentially malicious traffic even if only

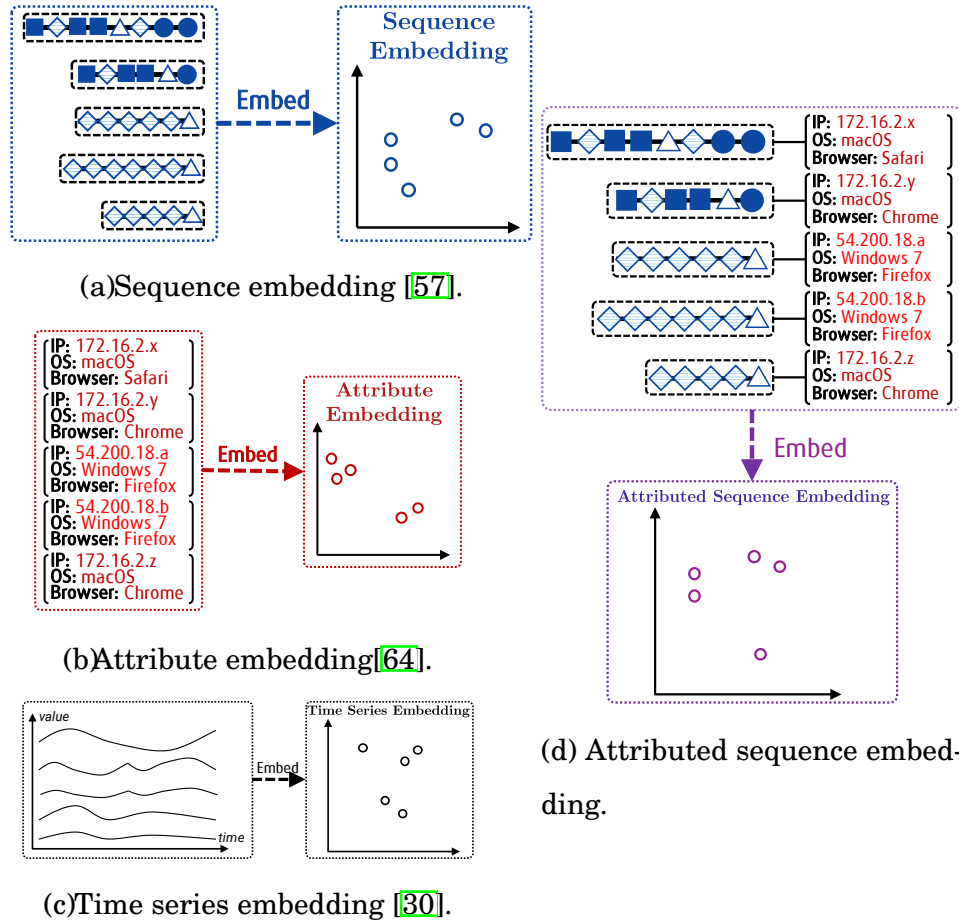


Figure 1.4: Comparison of different embedding problems.

one or a few examples per known intrusion type have been seen previously (as depicted in Figure 1.3).

1.3 State-of-the-Art

1.3.1 Attributed Sequence Embedding

Sequential data usually requires a careful design of its feature representation before being fed to a learning algorithm. One of the feature learning problems on sequential data is called sequence embedding [12, 57], where the goal is to transform a sequence into a fixed-length feature representation. Similarly, the attributed sequence embedding problem corresponds to transforming an attributed sequence into a fixed-length feature representation with continuous values.

In the sequence context, conventional methods only focus on sequential data [57, 40, 12, 27, 46] to learn the dependencies between items within variable-length sequences – neither support the attribute data nor learn the dependencies between attributes and sequences.

On the other hand, the recent applications on image datasets using multilayer deep neural networks [64, 1, 26, 56] focus on the problems of pattern recognition in fixed-size images – they neither support the variable-length sequences nor learn the attribute-sequence dependencies. Here in Figure 1.4, we demonstrate the difference between our

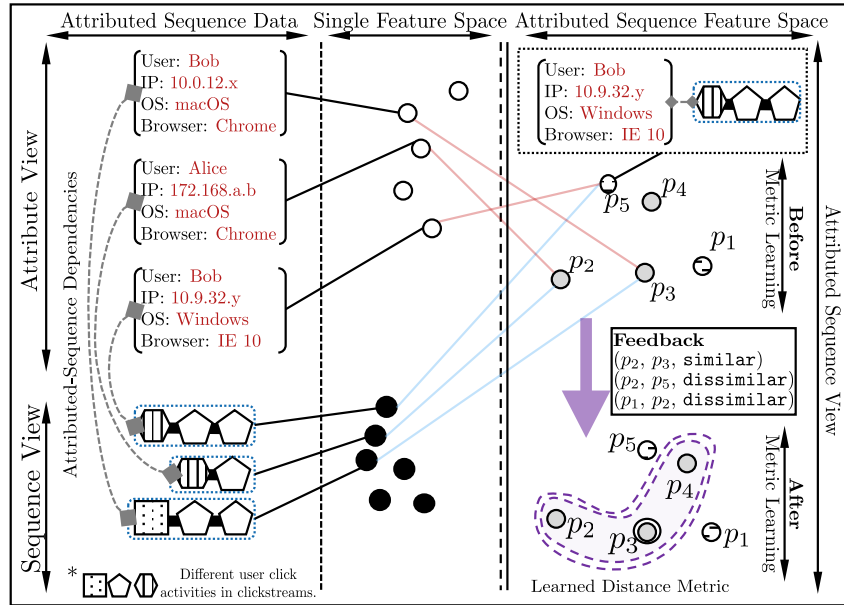


Figure 1.5: Distance metric learning on attributed sequences.

attributed sequence embedding problem and other embedding problems in the state-of-the-art.

1.3.2 Deep Metric Learning on Attributed Sequences

Conventional approaches on distance metric learning [70, 14, 33, 43] mainly focus on learning a Mahalanobis distance metric, which is equivalent to learning a *linear transformation* on data attributes. Recent research has extended distance metric learning to nonlinear settings [29, 26], where a nonlinear mapping function is first learned to project the instances into a new space, and then the final metric becomes the Euclidean distance metric in that space.

Deep metric learning has been the method of choice in practice for learning nonlinear mappings [63, 9, 29, 26]. Recent research on metric learning has explored sequential data [46], where we have structural information in the sequences, but no attributes are available. We use Figure 1.5 to depict the differences in distance metrics learned from the data.


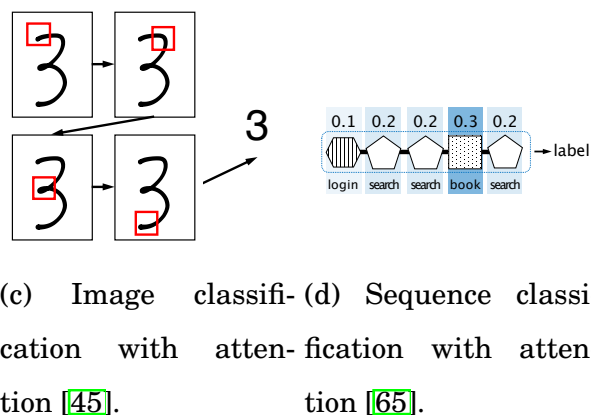



Figure 1.6: A comparison of related classification problems.

1.3.3 One-shot Learning on Attributed Sequences

With the capacity of deep learning for generalizing the training examples, one-shot learning has attracted more interests recently when only one training sample is available [54, 6, 59, 32]. Conventional approaches to one-shot learning focus on using feature vectors as input in the learning process [32, 63, 71], in which each instance is represented as a fixed-size vector (*e.g.*, images). However, neither sequential data nor attributed sequences has been studied in one-shot learning research.

1.3.4 Attention for Attributed Sequence Classification

In the areas of image processing, researchers designed a mechanism of only letting the model learn from certain “useful” and “informative” regions of image inputs [45]. This mechanism is called the attention network. Attention network has gained more research interest in recent years and has been generalized in various domains, *e.g.*, image captioning [72, 48] and generation [22], speech recognition [13], document classification [76]. However, attention model has yet been studied for attributed sequence classification. We use Figure 1.6 to depict different classification problems on attributed sequences.

1.4 Research Challenges

Here we summarize the research challenges of using attributed sequences in data mining applications:




Figure 1.7: The three types of dependencies in an attributed sequence: **item** dependencies, **attribute** dependencies and **attribute-sequence** dependencies.

1.4.1 Heterogeneous dependencies in attributed sequences

The bipartite structure of attributed sequences poses unique challenges in feature learning. Contrary to a straightforward thinking that attributes and sequences could first be learned separately then later re-joined, various dependencies arise in attributed sequences. There exist three types of dependencies in an attributed sequence: item dependencies, attribute dependencies and attribute-sequence dependencies. Thus, learning and capturing these *attribute-sequence dependencies* are criti-

cal for attributed sequence embedding. For example, in a web search, each user session is composed of a session profile as attributes (*e.g.*, Device Type, OS, *etc*) as well as a sequence of search keywords. One keyword may depend on previous search terms (*e.g.*, temperature following snow storm) and the keywords searched may depend on the device type (*e.g.*, Nearest restaurant on Cellphone). In the user behavior studies using web search histories, it would be impractical without considering such dependencies. We illustrate these three types of dependencies in Figure [1.7](#).

1.4.2 Lack of labeled data

With the continuously incoming volume of data and the high labor cost of manually labeling data, it is rare to find attributed sequences from many real-world applications with labels (*e.g.*, *fraud*, *normal*) attached. Without proper labels, it is challenging to learn an embedding function that is capable of transforming attributed sequences into compact feature representations while capturing the three types of dependencies.

1.4.3 Metric Learning on Attributed Sequences

Metric learning focuses on the problems of differentiating the instances based on the similarities/dissimilarities in user feedback. One fundamental problem in metric learning on attributed sequences lies in the sequential structure in each instance. Conventional approaches to distance metrics learning assume, explicitly or implicitly, that the data are represented as features vectors (*i.e.*, attributes) [70, 14, 33, 43]. However, in attributed sequence data, the sequence in each data record is not represented as a feature vector, but rather as a variable-length sequence of discrete items. The information is encoded in the ordering of the items. A distance metric on attributed sequences thus must be able to capture structural similarities and dissimilarities between sequences.

1.4.4 Generalize from Only One Sample

This problem is different from previous one-shot learning work, as we now need to extract feature vectors from not only the *attributes* but also the structural information from the *sequences* and the *dependencies* between attributes and sequences. The key difficulty in one-shot learning is to generalize beyond the single training example. It is more difficult to generalize from a more complex data type [11], such as attributed

sequence data, than from a simpler data type due to the larger search space and slower convergence.

1.4.5 Attention in Attributes and Sequences

Recent sequence learning research [76] has used neural attention models to improve the performance of sequence learning, such as document classification [76]. However, the attention mechanism focuses on learning the weight of certain time steps or sub-sequences in each sequential instance, without regards to its associated attributes. With information from the attributes, the weight of item or sub-sequence may be drastically different from the weight calculated by the attention mechanism using only sequences, which would consequently have different classification results.

1.5 Proposed Solutions

In this dissertation, we address the challenges listed in Section 1.4 by extending the state-of-the-art deep learning models under various problem settings.

1.5.1 Attributed Sequence Embedding

The goal of the first task is to generate fix-sized feature representations (*i.e.*, embeddings) of attributed sequences that can be used in downstream mining tasks (*e.g.*, clustering and outlier detection). Many of these machine learning tasks use vectors of real numbers as the inputs. The embeddings, as mappings from attributed sequences to fixed-size vectors of real numbers, are both important and valuable in these tasks. Because the embeddings are the vectors mapped from inputs, the similarities measured by distance functions in the embedding space can be viewed as the similarities between the original inputs. This task offers the following contributions:

- We study the problem of attributed sequence embedding, which corresponds to a natural generalization of many real-world applications. We formally define the data model of attributed sequences, which consists of a sequence of categorical items and a set of static attributes.
- We propose an innovative framework to exploit the dependencies among the attributed sequences. The NAS establishes the first attributed sequence embedding framework, which employs a three-

phase process to effectively produce feature representations for attributed sequences.

- We evaluate the NAS framework on various real-world datasets competing with state-of-the-art methods. We study the performance of NAS using clustering and outlier detection tasks. We also show that NAS is capable of producing feature representations in real-time. We also evaluate NAS using real-world case studies of user behaviors.

1.5.2 Incorporating Feedback of Attributed Sequences

Different from Task 1, now there is a limited amount of labeled feedback from domain experts in some real-world applications, such as *fraud detection* and *user behavior analysis*. The feedback is usually given as pairwise examples, where each feedback instance is composed of two attributed sequences and a label depicting whether they are similar or dissimilar. The feedback from domain experts is important in these applications since the feedback provides human insights of the datasets. For example, a domain expert can give feedback indicating a new online transaction is similar to a known fraud transaction. To address the above challenges, we propose a deep learning framework, called MLAS

(Metric Learning on Atttributed Sequences), to learn a distance metric that can effectively measure the dissimilarity between attributed sequences. Our MLAS framework includes three sub-networks: `AttNet` (an attribute network to encode the attribute information using nonlinear transformations), `SeqNet` (a sequence network to encode structural information using LSTM), and `MetricNet` (a metric network to produce the distance metric). We further designed three MLAS models: `balanced`, `AttNet-centric` and `SeqNet-centric`. We offer the following main contributions in this task:

- We are the first to formulate and then study the problem of deep metric learning of attributed sequences.
- We design three sub-networks: the `AttNet` to encode the attribute information, `SeqNet` to encode the structural information, and a metric network `MetricNet` to produce the distance metric. Together with these sub-networks, MLAS effectively learns the non-linear distance metric on attributed sequences.
- Our empirical results on real-world datasets demonstrate that our proposed MLAS framework significantly improves the performance of metrics learning on attributed sequences.

1.5.3 Classify Attributed Sequences in One-shot

To address the challenges in Section ??, we propose an end-to-end one-shot learning model, called OLAS, to accomplish one-shot learning for attributed sequences. The OLAS model includes two main components: a `CoreNet` to encode the information from attributes, sequences and their dependencies and a `PredictNet` to learn the similarities and differences between different attributed sequence classes. The proposed OLAS model is beyond a simple concatenation of `CoreNet` and `PredictNet`. Instead, they are interconnected within one network architecture and thus can be trained synchronously. Once the OLAS is trained, we can then use it to make predictions for not only the new data but also for entire previously unseen new classes. In this task, we offer the following core contributions:

- We formulate and analyze the problem of one-shot learning for attributed sequence classification.
- We develop a deep learning model that is capable of inferring class labels for attributed sequences based on one instance per class.

- We demonstrate that the OLAS network model trained on attributed sequences significantly improves the accuracy of label prediction compared to state-of-the-art.

1.5.4 Attention Model for Attributed Sequence Classification

When presented an image, humans are capable of only focusing on some regions of the image and grasping the information that is deemed useful. The attention models in deep learning are designed to fulfill similar goals. Recently, the attention models have generated a lot of research interests in various areas [45, 22, 13, 76]. For instance, *attention* has been employed to improve the performance of natural language processing tasks (e.g., document classification). Compared to treating every word and sentence equally in the document classification tasks without attention, the capability of enhanced processing of certain words or sentences using attention has demonstrated improvements in such tasks [76]. In this task, I design two attention models for attributed sequence classification. Specifically, this task offers the following contributions:

- We formulate the problem of attributed sequence classification.

- We design a deep learning framework, called AMAS, with two attention-based models to exploit the information from attributes and sequences.
- We demonstrate that the proposed models significantly improve the performance of attributed sequence classification using performance experiments and case studies.

1.6 Road Map

This rest of the dissertation is organized as follows. We introduce the attributed sequence data model in Chapter [2](#). We start from the first task of attributed sequence embedding in Chapter [3](#). Chapter [4](#) focuses on the task of deep metric learning on attributed sequences. In Chapters [5](#) and [6](#), we present the tasks of one-shot learning and the attention network for attributed sequence classification tasks, respectively. Related work is discussed in Chapter [7](#). We conclude the findings and discuss future works in Chapter [8](#).

Chapter 2

Background: Attributed

Sequence

Definition 1 (Sequence.) Given a set of r categorical items $\mathcal{I} = \{e_1, \dots, e_r\}$, the k -th sequence in the dataset $S_k = (\alpha_k^{(1)}, \dots, \alpha_k^{(l_k)})$ is an ordered list of items, where $\alpha_k^{(t)} \in \mathcal{I}, \forall t = 1, \dots, l_k$.

Different sequences can have a varying number of items. For example, the number of user click activities varies between different user sessions. The meanings of items are different in different datasets. For example, in user behavior analysis from clickstreams, each item represents one action in user's click stream (e.g., $\mathcal{I} = \{\text{search}, \text{select}\}$, where $r = 2$). Similarly in DNA sequencing, each item represents one canonical base (e.g., $\mathcal{I} = \{\text{A}, \text{T}, \text{G}, \text{C}\}$, where $r = 4$). There are dependencies between items in a sequence. Without loss of generality, we use the one-hot encoding of S_k , denoted as $\mathbf{S}_k = (\boldsymbol{\alpha}_k^{(1)}, \dots, \boldsymbol{\alpha}_k^{(l_k)}) \in \mathbb{R}^{l_k \times r}$ where each item $\boldsymbol{\alpha}_k^{(t)} \in \mathbb{R}^r$ in \mathbf{S}_k is a one-hot vector corresponding to the original item $\alpha_k^{(t)}$ in the sequence S_k .

In prior work [21], one common preprocessing step is to zero-pad each sequence to the longest sequence in the dataset and then to one-hot encode it. Without loss of generality, we denote the length of the longest sequence as T . For each k -th sequence S_k in the dataset, we denote its equivalent one-hot encoded sequence as $\mathbf{S}_k = (\boldsymbol{\alpha}_k^{(1)}, \dots, \boldsymbol{\alpha}_k^{(T_k)})$, where

$\alpha_k^{(t)} \in \mathbb{R}^r$ corresponds to a vector represents the item $\alpha_k^{(t)}$ with the l -th entry in $\alpha_k^{(t)}$ is “1” and all other entries are zeros if $\alpha_k^{(t)} = e_l, e_l \in \mathcal{I}$.

Definition 2 (Attributes.) The attribute values are concatenated into a vector $\mathbf{x}_k \in \mathbb{R}^u$, where u is the number of attributes in \mathbf{x}_k .

The value of each attribute can be either categorical or numerical. u is considered as a constant within any given datasets. For example, in a dataset where each instance has two attributes “IP” and “OS”, $u = 2$.

Definition 3 (Attributed Sequence.) Given a vector of attribute values \mathbf{x}_k and a sequence \mathbf{S}_k , an attributed sequence $J_k = (\mathbf{x}_k, \mathbf{S}_k)$ is an ordered pair of the attribute value vector \mathbf{x}_k and the sequence \mathbf{S}_k .

Common practices of distance metric learning involve pairwise examples as the training dataset [60, 61, 77, 33, 26, 46]. We thus define the feedback as a collection of similar (or dissimilar) attributed sequence pairs.

Definition 4 (Similar and Dissimilar Feedback) Let $\mathcal{J} = \{J_1, \dots, J_n\}$ be a set of n attributed sequences. A feedback is a triplet (p_i, p_j, ℓ_{ij}) consisting of two attributed sequences $p_i, p_j \in \mathcal{J}$ and a label $\ell_{ij} \in \{0, 1\}$ indicating whether p_i and p_j are similar ($\ell_{ij} = 0$) or dissimilar ($\ell_{ij} = 1$).

We define a similar feedback set $\mathcal{S} = \{(p_i, p_j, \ell_{ij}) | \ell_{ij} = 0\}$ and a dissimilar feedback set $\mathcal{D} = \{(p_i, p_j, \ell_{ij}) | \ell_{ij} = 1\}$.

Chapter 3

Unsupervised Attributed

Sequence Embedding

3.1 Problem Definition

Using the previously defined notations in Chapter 2, we formulate our attributed sequence embedding problem as follows.

Definition 5 (Attributed Sequence Embedding.) Given a dataset of attributed sequences $\mathcal{J} = \{J_1, \dots, J_n\}$, the problem of attributed sequence embedding is to find a function Φ parameterized by θ that produces feature representations for J_k in the form of vectors. The problem is formulated as

$$\text{minimize } \sum_{k=1}^n \sum_{t=1}^{l_k} -\log \Pr_{\phi} \left(\alpha_k^{(t)} | \beta_k^{(t)}, \mathbf{x}_k \right) \quad (3.1)$$

where $\beta_k^{(t)} = (\alpha_k^{(t-1)}, \dots, \alpha_k^{(1)})$, $\forall t = 2, \dots, l_k$ represents the items prior to $\alpha_k^{(t)}$ in the sequence.

Our problem can be interpreted as: we want to minimize the prediction error of the $\alpha_k^{(t)}$ in each attributed sequence given the attribute values \mathbf{x}_k and all the items prior to $\alpha_k^{(t)}$.




Figure 3.1: The process flow of the NAS framework from attributed sequences to the final feature representations is displayed. For an attributed sequence, the attribute feature representation from the attribute network is shared with the sequence network. In the end, the hidden layer states of the sequence network are saved as the feature representation for this attributed sequence.

3.2 The NAS Framework

In this section, we introduce two main components in our proposed NAS framework, *i.e.*, *attribute network* and *sequence network*. We illustrate how the two networks are integrated to serve the purpose of embedding attributed sequences in Figure 3.1.

In many real-world applications, attributes are often available before the sequences. For example, in the online ticketing system, the attributes that depict user profile are recorded when a user starts a ticket booking session. However, the sequence of one’s actions on the webpage would not be available until such booking session ends. Our design follows the same paradigm, that is, the attributes are available before the respective sequence of items in each attributed sequence. In addition,

the lengths of sequence in many real-world applications are short. For example, one may only need tens of actions to finish a ticket booking online.

3.2.1 Attribute Network

Fully connected neural network [38] is capable of modeling the dependencies of the inputs and at the same time reduce the dimensionality. It has been widely used [38, 39, 53] for unsupervised data representations learning, including tasks such as dimensionality reduction and generative data modeling. With the high-dimensional sparse input attribute values $\mathbf{x}_k \in \mathbb{R}^u$, it is ideal to use such a network to learn the attribute dependencies. We design our attribute network as

$$\begin{aligned}
 \mathbf{V}_k^{(1)} &= \rho \left(\mathbf{W}_A^{(1)} \mathbf{x}_k + \mathbf{b}_A^{(1)} \right) \\
 &\vdots \\
 \mathbf{V}_k^{(M)} &= \rho \left(\mathbf{W}_A^{(M)} \mathbf{V}_k^{(M-1)} + \mathbf{b}_A^{(M)} \right) \\
 \mathbf{V}_k^{(M+1)} &= \sigma \left(\mathbf{W}_A^{(M+1)} \mathbf{V}_k^{(M)} + \mathbf{b}_A^{(M+1)} \right) \\
 &\vdots \\
 \widehat{\mathbf{x}}_k &= \sigma \left(\mathbf{W}_A^{(2M)} \mathbf{V}_k^{(2M-1)} + \mathbf{b}_A^{(2M)} \right)
 \end{aligned} \tag{3.2}$$

where ρ is the encoder activation function and σ is the decoder activation function. In this attribute network, we use ReLU activation function [47],

defined as

$$\rho(z) = \max(0, z) \quad (3.3)$$

and sigmoid activation function, defined as

$$\sigma(z) = \frac{1}{1 + e^{-z}} \quad (3.4)$$

We denote the parameter set of attribute network as $\phi_A = \{\mathbf{W}_A, \mathbf{b}_A\}$, where $\mathbf{W}_A = (\mathbf{W}_A^{(1)}, \dots, \mathbf{W}_A^{(2M)})$ and $\mathbf{b}_A = (\mathbf{b}_A^{(1)}, \dots, \mathbf{b}_A^{(2M)})$.

The attribute network with $2M$ layers has two components, *i.e.*, the encoder, and the decoder. The encoder is composed of the first M layers, and the next M layers work as the decoder. With d_M hidden units in the M -th layer, the input attribute vector $\mathbf{x}_k \in \mathbb{R}^u$ is first transformed to $\mathbf{V}_k^{(M)} \in \mathbb{R}^{d_M}$, $d_M \ll u$ by the encoder. Then the decoder attempts to reconstruct the input and produce the reconstruction result $\widehat{\mathbf{x}}_k \in \mathbb{R}^u$. An ideal attribute network should be able to reconstruct the input from the $\mathbf{V}_k^{(M)}$. The smallest attribute network is built with $M = 1$, where there are one encoder and one decoder.

3.2.2 Sequence Network

The proposed sequence network is a variation of the long short-term memory model (LSTM). The sequence network takes advantage of LSTM to learn the dependencies between items in sequences. Sequence network also accepts the feature representations from the attribute network to affect the learning process and thus learn the attribute-sequence dependencies. We define our sequence network as

$$\begin{aligned}
\mathbf{i}_k^{(t)} &= \sigma \left(\mathbf{W}_i \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_i \mathbf{h}_k^{(t-1)} + \mathbf{b}_i \right) \\
\mathbf{f}_k^{(t)} &= \sigma \left(\mathbf{W}_f \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_f \mathbf{h}_k^{(t-1)} + \mathbf{b}_f \right) \\
\mathbf{o}_k^{(t)} &= \sigma \left(\mathbf{W}_o \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_o \mathbf{h}_k^{(t-1)} + \mathbf{b}_o \right) \\
\mathbf{g}_k^{(t)} &= \sigma \left(\mathbf{W}_g \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_g \mathbf{h}_k^{(t-1)} + \mathbf{b}_g \right) \\
\mathbf{c}_k^{(t)} &= \mathbf{f}_k^{(t)} \odot \mathbf{c}_k^{(t-1)} + \mathbf{i}_k^{(t)} \odot \mathbf{g}_k^{(t)} \\
\mathbf{h}_k^{(t)} &= \mathbf{o}_k^{(t)} \odot \tanh \left(\mathbf{c}_k^{(t)} \right) + \mathbf{1}(t=1) \odot \mathbf{V}_k^{(M)} \\
\mathbf{y}_k^{(t)} &= \delta \left(\mathbf{W}_y \mathbf{h}_k^{(t)} + \mathbf{b}_y \right)
\end{aligned} \tag{3.5}$$

where σ is a sigmoid activation function, $\mathbf{i}_k^{(t)}$, $\mathbf{f}_k^{(t)}$, $\mathbf{o}_k^{(t)}$ and $\mathbf{g}_k^{(t)}$ are the internal gates. With d hidden units in the sequence network, $\mathbf{c}_k^{(t)} \in \mathbb{R}^d$, $\mathbf{h}_k^{(t)} \in \mathbb{R}^d$ are the cell states and hidden states of the sequence network. $\mathbf{y}_k^{(t)} \in \mathbb{R}^r$ is the predicted next item in sequence computed using

softmax function. With the softmax activation function defined as

$$\delta(\mathbf{z}) = \frac{\exp(z_j)}{\sum_{i=1}^r \exp(z_i)}, \forall j = 1, \dots, r \quad (3.6)$$

where the i -th element in \mathbf{z} is denoted as z_i , the $y_k^{(t)}$ can be interpreted as the probability distribution over r items.

We integrate the attribute network by conditioning the hidden states in the sequence network at the first time step (denoted as $\mathbb{1}(t = 1) \odot \mathbf{V}_k^{(M)}$). To do this, the attribute network and sequence network should have the same number of hidden units, *i.e.*, $d_M = d$. After processing the last time step for an attributed sequence \mathbf{S}_k , the cell state of sequence network, namely $\mathbf{c}_k^{(l_k)}$, is used as the feature representation of \mathbf{S}_k . We denote the set of parameters in sequence network as $\phi_S = \{\mathbf{W}_S, \mathbf{U}_S, \mathbf{b}_S, \mathbf{W}_y, \mathbf{b}_y\}$, where $\mathbf{W}_S = (\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_g)$, $\mathbf{U}_S = (\mathbf{U}_i, \mathbf{U}_f, \mathbf{U}_o, \mathbf{U}_g)$ and $\mathbf{b}_S = (\mathbf{b}_i, \mathbf{b}_f, \mathbf{b}_o, \mathbf{b}_g)$.

3.2.3 Attributed Sequence Learning

The attributed sequence embeddings would not be useful for downstream mining tasks if the embedding space is detached from the inputs. Thus, we choose two different learning objective functions for the attribute net-

work and sequence network targeting at the different characteristics of attribute and sequence data.

Attribute network aims at minimizing the differences between the input and reconstructed attribute values. The learning objective function of attribute network is defined as:

$$L_A = \|\mathbf{x}_k - \widehat{\mathbf{x}}_k\|_2^2 \quad (3.7)$$

Sequence network aims at minimizing log likelihood of the incorrect prediction of the next item at each time step. Thus, the sequence network learning objective function can be formulated using categorical cross-entropy as:

$$L_S = - \sum_{t=1}^{l_k} \alpha_k^{(t)} \log \mathbf{y}_k^{(t)} \quad (3.8)$$

Also, the learning processes are composed of a number of iterations, and the parameters are updated during each iteration based on the gradient computed. Without loss of generality, we denote L_A^τ and L_B^τ as the τ -th iteration of attribute network and sequence network, respectively. We further denote the maximum numbers of iterations for attribute network and sequence network as \mathcal{T}_A and \mathcal{T}_B . \mathcal{T}_A and \mathcal{T}_B may not be equal as

Algorithm 1 NAS Learning

INPUT: $\mathcal{J} = \{J_1, \dots, J_n\}$, attribute network layers $2M$, learning rate γ , iteration number of attribute network \mathcal{T}_A and sequence network \mathcal{T}_S and convergence error ε_A and ε_S .

OUTPUT: Parameter sets ϕ_A, ϕ_S .

```

1: Initialize  $\phi_A$  and  $\phi_S$  using uniform distribution.
2: for each  $J_k \in \mathcal{J}, k = 1, \dots, n$  do
3:   for each  $\tau = 1, \dots, \mathcal{T}_A$  do
4:     for each  $m = 1, \dots, 2M$  do
5:       Compute forward propagation.
6:     end for
7:     for each  $m = 2M, \dots, 1$  do
8:       Compute the gradient of layer  $m$ .
9:     end for
10:    for each  $m = 1, \dots, 2M$  do
11:      Update the parameter set  $\phi_A$ 
12:    end for
13:    Calculate  $L_A^\tau$ 
14:    if  $\tau > 1$  and  $|L_A^\tau - L_A^{\tau-1}| < \varepsilon_A$  then
15:      Stop iterating.
16:    end if
17:  end for
18:  for each  $\tau = 1, \dots, \mathcal{T}_S$  do
19:    for each  $t = 1, \dots, l_k$  do
20:      Compute forward propagation and get  $y_k^t$ .
21:      Compute the gradients of sequence network  $\Phi_S$ .
22:      Update the parameter set  $\phi_S$ .
23:    end for
24:    Calculate  $L_S^\tau$ 
25:    if  $\tau > 1$  and  $|L_S^\tau - L_S^{\tau-1}| < \varepsilon_S$  then
26:      Stop iterating.
27:    end if
28:  end for
29: end for

```

the number of iterations needed for the attribute network, and sequence network may not be the same. Algorithm [1](#) summarizes the learning

paradigm under our proposed NAS framework. After the attributed sequence learning process, we use the parameters in the attribute network and sequence network to embed each attributed sequence.

In this section, we evaluate NAS framework using real-world application logs from Amadeus and public datasets from Wikispeedia [67, 68]. We evaluate the quality of embeddings generated by different methods by measuring the performance of outlier detection and clustering algorithms using different embeddings. Outlier detection and clustering tasks are frequently being used for many applications, such as fraud detection and user behavior analysis. We also include three methods not using neural networks in outlier detection tasks. We also study four case studies in the security management system in Amadeus to demonstrate the embeddings produced by NAS is useful for real-world applications. Lastly, we show that the NAS framework is capable of embedding attributed sequences in real-time.

3.3 Experimental Setup

3.3.1 Data Collection

We use four datasets in our experiments: two from Amadeus application log files and two from Wikispeedia^[1]. The numbers of attributed sequences in all four datasets vary between $\sim 58k$ and $\sim 106k$.

- **AMS-A:** We extract $\sim 58k$ instances from log files of an Amadeus internal application. Each record is composed of a profile containing information ranging from system configurations to office name, and a sequence of functions invoked by click activities on the web interface. There are 288 distinct functions, 57,270 distinct profiles in this dataset. The average length of the sequences is 18.
- **AMS-B:** We use $\sim 106k$ instances from Amadeus internal application log files with 573 distinct functions and 106,671 distinct profiles. The average length of the sequences is 22.
- **Wiki-A:** This dataset is sampled from Wikispeedia dataset^[2]. Wikispeedia dataset originated from a human-computation game, called Wikispeedia [68]. In this game, each user is given two pages (*i.e.*,

¹Personal identity information are not collected for privacy reasons.

²Download link: <http://goo.gl/8Z9h9f>

source, and destination) from a subset of Wikipedia pages and asked to navigate from the source to the destination page. We use a subset of $\sim 2k$ paths from Wikispeedia with the average length of the path as 6. We also extract 11 sequence context (*e.g.*, the category of the source page, average time spent on each page) as attributes.

- `Wiki-B`: This dataset is also sampled from Wikispeedia dataset. In `Wiki-B`, we use $\sim 1.5k$ paths from Wikispeedia with the average length of the path as 8. We also extract 11 sequence context as attributes.

3.3.2 Compared Methods

To study NAS performance on attributed sequences in real-world applications, we use the compared methods in Table [3.1](#) in outlier detection and clustering tasks.

- `LEN` [\[1\]](#): The attributes are encoded and directly used in the mining algorithm.
- `MCC` [\[5\]](#): `MCC` uses the sequence component of attributed sequence as input and produces log likelihood for each sequence.

Table 3.1: Compared methods.

Methods	Data Used	Short Descriptions	Related Publication
LEN	Attributes	The encoded attributes is used.	[1]
MCC	Sequences	A markov chain based method.	[5]
ATR	Attributes	An autoencoder for attribute data.	[64]
SEQ	Sequences	Sequence embedding using LSTM.	[57]
EML	Attributes Sequences	Aggregation of the scores of LEN and MCC methods.	[74]
CSA	Attributes Sequences	Concatenation of embeddings generated by ATR and SEQ methods.	[50]
NAS	Attributes Sequences	Attribute network is used to condition sequence network.	This Work

- SEQ [57]: Only the sequence inputs are used by an LSTM to generate fixed-length embeddings.
- ATR [64]: A fully connected neural network with two layers is applied on only attributes to generate embeddings.
- EML [74]: The scores from MCC and LEN are aggregated.
- CSA [50]: The attribute embedding and the sequence embedding are first generated by ATR and SEQ, respectively. Then the two embeddings are concatenated together.

-
- NAS: Our proposed NAS framework using both attributes and sequences to generate embeddings.

3.3.3 Network Parameters

Following the previous work in [20], we initialize weight matrices W_A and W_S using the uniform distribution. The recurrent matrix U_S is initialized using the orthogonal matrix as suggested by [55]. All the bias vectors are initialized with zero vector 0 . We use stochastic gradient descent as optimizer with the learning rate of 0.01. We use a two-layer encoder-decoder stack as our attribute network.

3.4 Outlier Detection Tasks

We first use outlier detection tasks to evaluate the quality of embeddings produced by NAS. We select k -NN outlier detection algorithm as it has only one important parameter (*i.e.*, the k value). We use ROC AUC as the metric in this set of experiments.

For each of the `AMS-A` and `AMS-B` datasets, we ask domain experts to select two users as inlier and outlier. These two users have completely different behaviors (*i.e.*, sequences) and metadata (*i.e.*, attributes). The

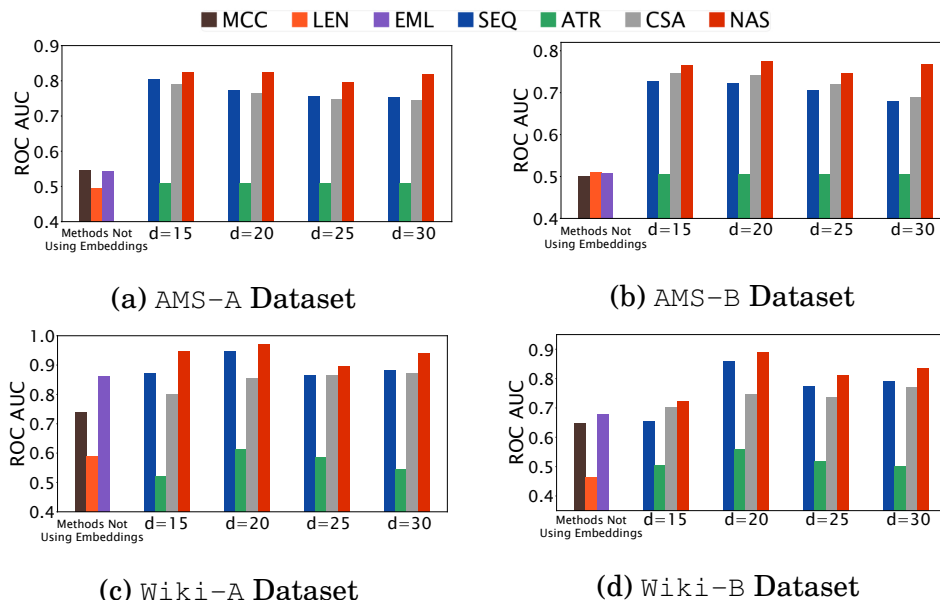


Figure 3.2: The performance of k -NN outlier detection ($k = 5$). The *methods not using embeddings* are placed on the left. We vary the number of dimensions on the right. The higher score is better. We observe that the combinations of k -NN and NAS embeddings have the best performance on the four datasets. percentages of outliers in AMS-A and AMS-B are 1.5% and 2.5% of all attributed sequences, respectively. For the Wiki-A and Wiki-B datasets, each path is labeled based on the category of the source page. Similarly to the previous two datasets, we select paths with two labels as inliers and outliers where the percentage of outlier paths is 2%. The feature used to label paths is excluded from the learning and embedding process.

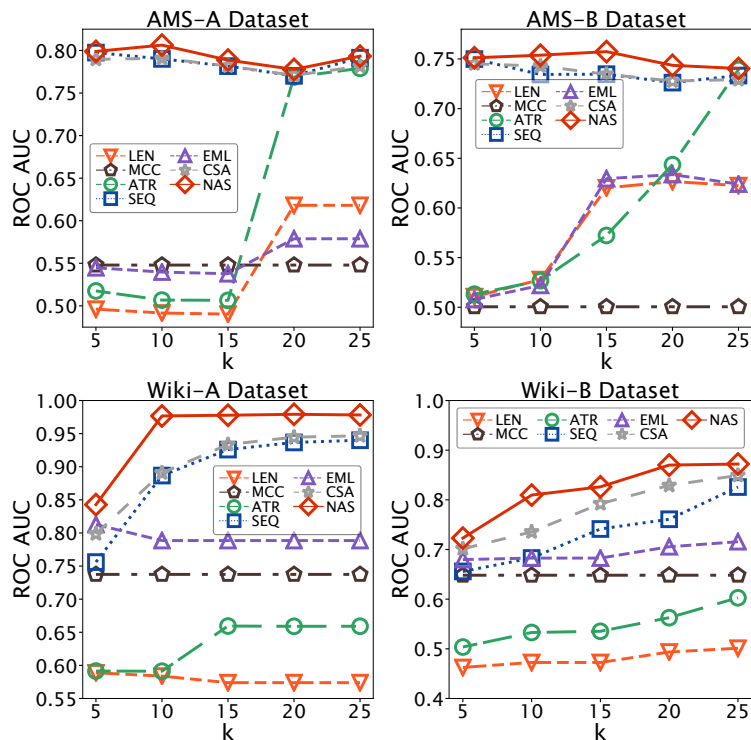


Figure 3.3: Parameter sensitivity to different k values. It is shown that the embeddings generated by NAS always have the best performance under different k values.

3.4.1 Performance

The goal of this set of experiments is to demonstrate the performance of outlier detection using all our compared methods. Each method is trained with all the instances. For SEQ, ATR and NAS, the number of learning epochs is set to 10 and we vary the number of embedding dimensions d from 15 to 30. We set $k = 5$ for outlier detection tasks for LEN, SEQ, ATR, CSA and NAS. Choosing the *optimal* k value in the out-

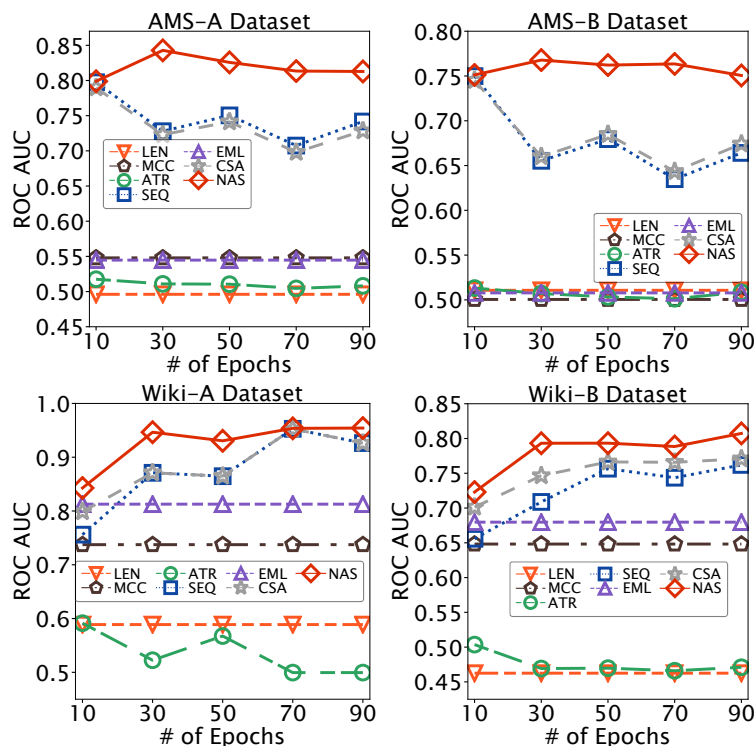


Figure 3.4: Performance comparisons using outlier detection tasks. The embeddings generated by NAS can always achieve the best performance compared to baseline methods when the number of training epochs increases.

lier detection tasks is beyond the scope of this work, thus we omit its discussions. We summarize the performance results in Figure [3.2](#).

Analysis. We find that the results based on the embeddings generated by NAS are superior to other methods. That is, NAS maximally outperforms other state-of-the-art algorithms by 32.9%, 27.5%, 44.8% and 48% on AMS-A, AMS-B, Wiki-A and Wiki-B datasets, respectively.

It is worth mentioning that NAS outperforms a similar baseline method CSA by incorporating the information of attribute-sequence dependencies.

3.4.2 Parameter Study

There are two *key* parameters in our evaluations, *i.e.*, k value for the k -NN algorithm and the learning epochs.

In Figure [3.3](#), we first show that the embeddings (dimension $d = 15$) generated by our NAS assist k -NN outlier detection algorithm to achieve superior performance under a wide range of k values ($k = 5, 10, 15, 20, 25$). We omit the detailed discussions of selecting the optimal k values due to the scope of this work.

Next, we evaluate the performance changes as the number of training epochs increases. We do not use the early stopping in this set of experiments as we want to fix the number of training epochs. Figure [3.4](#) presents the results when we fix $k = 5, d = 15$ and vary the number of epochs in the learning phase. We notice that compared to its competitor, the embeddings generated by NAS can achieve a higher AUC even with a relatively fewer number of learning epochs. Compared to other neural network-based methods (*i.e.*, SEQ, ATR and CSA), NAS have

a more stable performance. The NAS performance gain is not due to the advantage of using both attributes and sequences, but because of taking the various dependencies into account, as the other two competitors (*i.e.*, CSA and EML) also utilize the information from both attributes and sequences.

3.5 Clustering Tasks

We use clustering tasks to evaluate the quality of the embeddings generated by compared methods. We use the HDBSCAN clustering algorithm from [8] since the results remain stable over different runs, which is a fair and ideal choice for studying the performance of representation learning algorithms. We use normalized mutual information (NMI) to measure the quality of the clustering. The highest NMI score is 1.

3.5.1 Performance

In Figure 3.5 we report the result when the minimum cluster size is fixed to 640, and we vary the number of embedding dimensionality d from 5 to 35. NAS is capable of taking advantage of using the information not only from both attributes and sequences but also the attributed-sequence de-

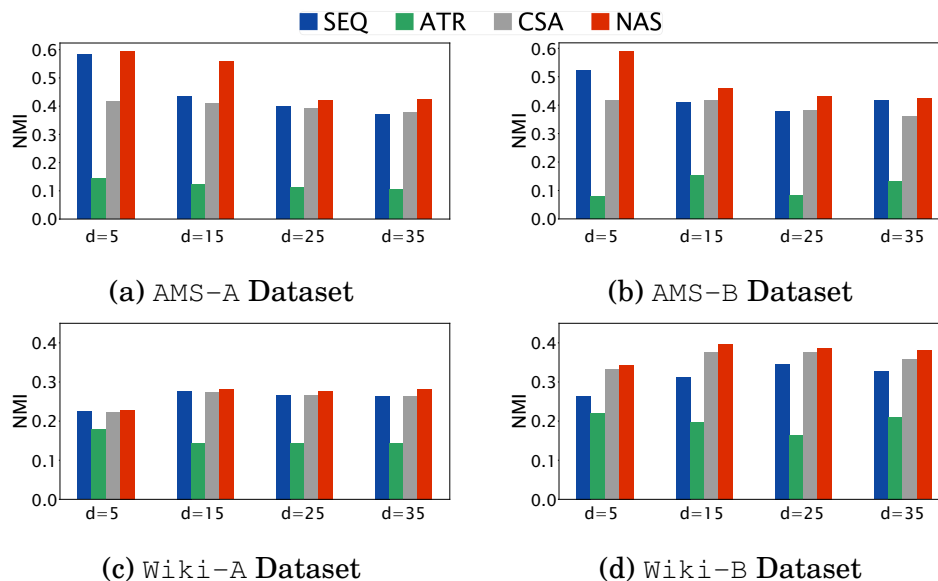


Figure 3.5: Performance of HDBSCAN clustering algorithm with the minimum cluster size of 640. The higher score is better. By using the embeddings produced by NAS, the clustering algorithm performs better than using embeddings produced by other baseline methods.

dependencies when generating the embeddings. Specially, when compared to the CSA method where the information from both attributes and sequences is used but without taking advantage of the attribute-sequence dependencies, NAS outperforms CSA by 24.7%, 20.6%, 3.8%, 4.3% on average on AMS-A, AMS-B, Wiki-A and Wiki-B datasets, respectively.

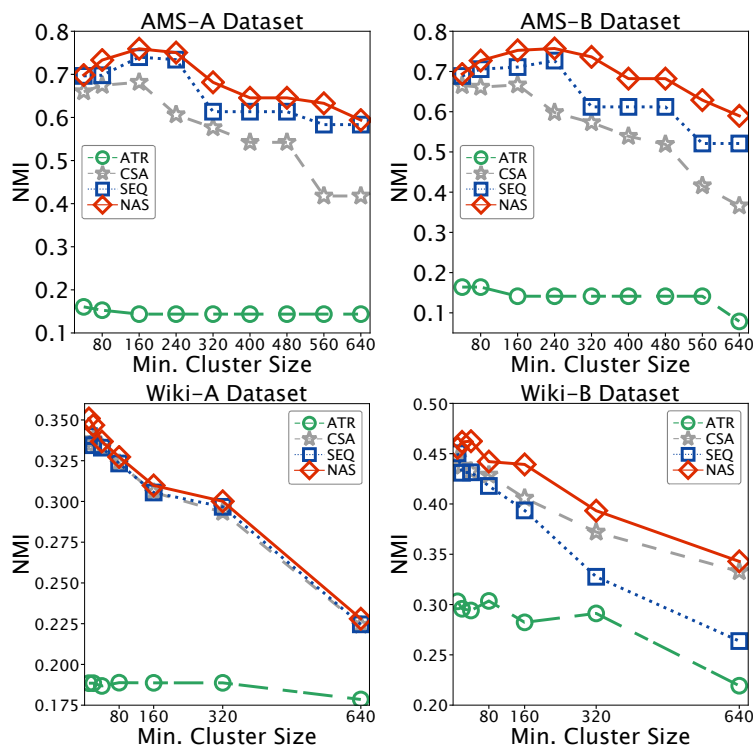


Figure 3.6: Clustering algorithm using the embeddings generated by NAS can achieve better performance than baselines under different minimum cluster size parameter.

3.5.2 Parameter Study

In this set of experiments, we investigate the performance of NAS under different parameter settings. We first vary the minimum size of clusters for all three datasets. For each dataset, we evaluate on the same range of the minimum cluster size. Figure 3.6 depicts the NMI of the HDBSCAN clustering algorithm when the dimensionality of embedding is fixed to 5,

and the number of learning epochs is set to 10. We observe that NAS is always achieving the highest NMI among the four methods.

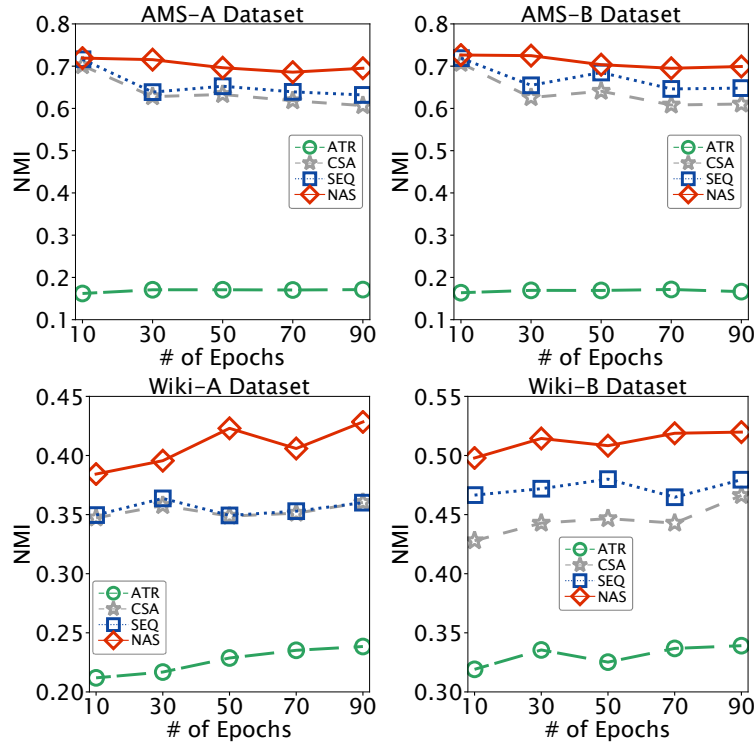


Figure 3.7: The embeddings generated by NAS have better performance than the baseline methods across different numbers of training epochs.

In Figure [3.7](#) we choose a fixed minimum cluster size of 40 and vary the number of learning epochs from 10 to 90. We observe that NAS is superior to its competitors across all parameter settings.

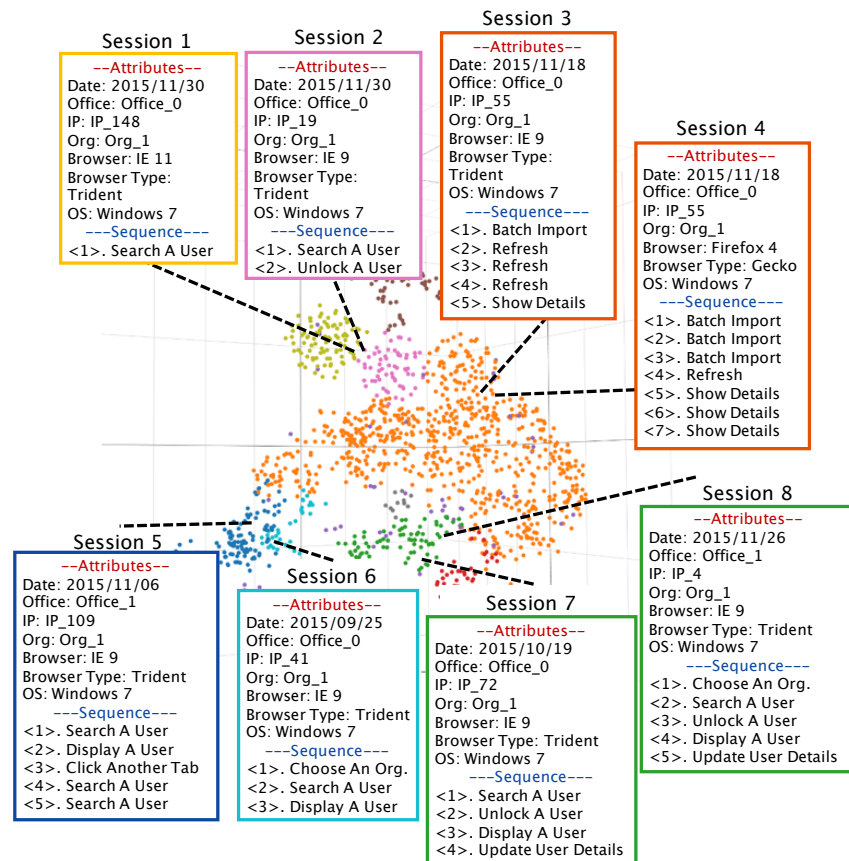


Figure 3.8: Clustering results for case studies. Each color represents one cluster. Only major attributes are shown here. Actual attribute values and item names in sequences are replaced for privacy reasons. (Better view in color.)

3.5.3 Case Study: Security Management System

In this case study, we build HDBSCAN clusters on the AMS-A dataset. Each point in Figure 3.8 is a user session in AMS-A dataset. We apply t-SNE [41] on the embeddings to generate a 3D plot with each color for one cluster. We ask domain experts to closely examine the similarities

and differences between the attributed sequences in each cluster. Based on eight user sessions in Figure [3.8](#), we summarize four case studies as follows:

- **Case Study 1.** The *similarities* and *differences* between the two “nearby” points from the *same* cluster. Sessions 3 and 4 share the same set of actions, and the order of each kind of action is the same. Further, they both share the same attributes. In a more complex case of Sessions 7 and 8, although they have different sets of actions (*i.e.*, Session 8 has “Choose An Org.” which Session 7 does not) and not all attributes are the same, they are clustered together since *both* of them are “high-level administrative” sessions.
- **Case Study 2.** *Differences* between two “nearby” points that belong to *different* clusters. In Sessions 1 and 2, although they share similar attributes (*i.e.*, “IE”, “Windows 7”, “Office_0”, etc.) and only one action is different, the Session 2 belongs to another cluster since that action (“Unlock A User”) changes the type of the session from “routine” to “administrative”.
- **Case Study 3.** *Similarities* between two “nearby” points that belong to *different* clusters. There are a number of differences be-

tween Sessions 5 and 6, which cause them to be distant from each other, such as Sessions 5 and 6 belong to different offices, have different IP addresses and have different sets of user actions in the sequences. The similarities between them are obvious: first, they both share a number of similar attributes (*i.e.*, “browser”, “OS”, and “organization”); second, the majority of actions remains the same, and they share a subsequence of actions (*i.e.*, “Search A User” and “Display A User”). Thus, although they belong to different clusters, the distance between Sessions 5 and 6 is smaller compared to sessions from other clusters.

- **Case Study 4.** A *global* view of the eight clusters. The NAS is capable of differentiating user sessions despite that the 8 user sessions have some similar or identical attribute values (*e.g.*, six of them are from the same office, all of them are using Windows 7), and there are common user actions shared by user sessions.

The above four case studies conclude that the clustering results based on the embeddings generated by NAS can be easily explained in real-world cases.

3.6 Scalability

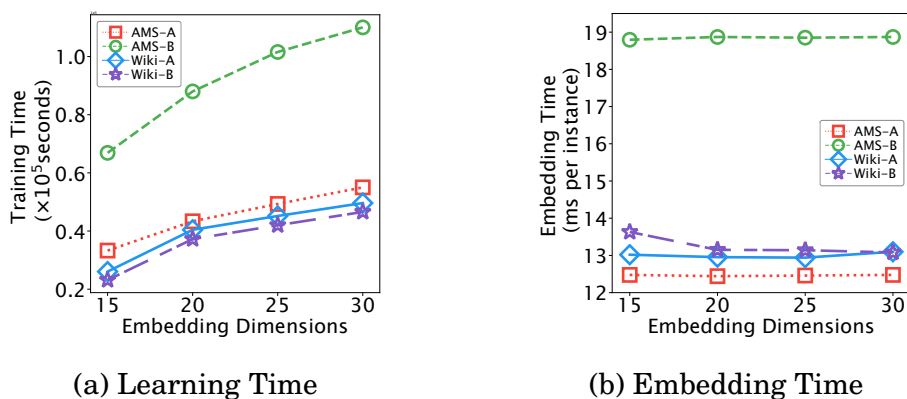


Figure 3.9: Runtime scalability

In this set of experiments, we evaluate the learning time and embedding time of the proposed NAS framework. We implement the NAS framework using Theano 0.8 [3] on Ubuntu 14.04. We conduct our experiments on a machine with 24 E5-2690 v2 cores and 256GB memory. The I/O time for reading datasets into memory and writing embeddings to the disk is excluded. Each setting corresponds to 10 runs times and averaged.

The embeddings of NAS come from the cell states in each artificial neuron. An increase in the embedding dimensions increases the number of artificial neurons in the model, which results in a larger model with more parameters that needs longer time to train. However, the model

learning process is designed to be an offline process. Thus, it does not interfere with the real-time attributed sequence embedding process. Although the learning time increases as the number of embedding dimensions increases, the embedding time per attributed sequence remains at the millisecond level. Thus, NAS is capable of transforming one attributed sequence into one embedding in real-time that is sufficient for real-world data mining tasks.

Chapter 4

Deep Metric Learning on Attributed Sequences

4.1 Problem Definition

Given a nonlinear transformation function Θ and two attributed sequences p_i and p_j as inputs, deep metric learning approaches [26] often apply the Mahalanobis distance function to the d -dimensional outputs of function Θ as

$$D_{\Theta}(p_i, p_j) = \sqrt{(\Theta(p_i) - \Theta(p_j))^{\top} \Lambda (\Theta(p_i) - \Theta(p_j))} \quad (4.1)$$

where $\Lambda \in \mathbb{R}^{d \times d}$ is a symmetric, semi-definite, and positive matrix. When $\Lambda = \mathbf{I}$, Eq. 5.3 is transformed to Euclidean distance [70] as:

$$D_{\Theta}(p_i, p_j) = \|\Theta(p_i) - \Theta(p_j)\|_2. \quad (4.2)$$

Given feedback sets \mathcal{S} and \mathcal{D} of attributed sequences as per Def. 4 and a distance function D_{Θ} as per Eq. 5.5, the goal of deep metric learning on attributed sequences is to find the transformation function $\Theta : (\mathbb{R}^u, \mathbb{R}^{T \times r}) \mapsto \mathbb{R}^d$ with a set of parameters θ that is capable of mapping the attributed sequence inputs to a space that the distances between each pair of attributed sequences in the similar feedback set \mathcal{S} are minimized while increasing the distances between attributed sequence pairs

in the dissimilar feedback set \mathcal{D} . Inspired by [70], we adopt the learning objective as

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \sum_{(p_i, p_j, \ell_{ij}) \in \mathcal{S}} D_{\theta}(p_i, p_j) \\ & \text{s.t.} && \sum_{(p_i, p_j, \ell_{ij}) \in \mathcal{D}} D_{\theta}(p_i, p_j) \geq g \end{aligned} \tag{4.3}$$

where g is a group-based margin parameter that stipulates the distance between two attributed sequences from dissimilar feedback set should be larger than g . This prevents the dataset from being reduced to a single point [70].

4.2 The MLAS Network

In this section, we first design two distinct networks, called `AttNet` and `SeqNet`, to learn the attribute data and sequence data, respectively. Next, we present several types of `FusionNet` designs to integrate the two networks. Lastly, the existing network composition is augmented by `MetricNet` to employ user feedback into the learning process.

4.2.1 AttNet and SeqNet

AttNet, designed to learn the relationships within attribute data, utilizes a fully connected neural network with multiple layers of nonlinear transformations. In particular, for an AttNet with M layers, we denote the weight and bias parameters of the m -th layer as $\mathbf{W}_A^{(m)}$ and $\mathbf{b}_A^{(m)}$, $\forall m = 1, \dots, M$. Given an attribute vector $\mathbf{x}_k \in \mathbb{R}^u$ as the input, with d_m hidden units in the m -th layer of AttNet, the corresponding output is $\mathbf{V}_k^{(m)} \in \mathbb{R}^{d_m}$, $\forall m = 1, \dots, M$. The structure of AttNet is designed as

$$\begin{aligned}
 \mathbf{V}_k^{(1)} &= \delta \left(\mathbf{W}_A^{(1)} \mathbf{x}_k + \mathbf{b}_A^{(1)} \right) \\
 \mathbf{V}_k^{(2)} &= \delta \left(\mathbf{W}_A^{(2)} \mathbf{V}_k^{(1)} + \mathbf{b}_A^{(2)} \right) \\
 &\vdots \\
 \mathbf{V}_k^{(M)} &= \delta \left(\mathbf{W}_A^{(M)} \mathbf{V}_k^{(M-1)} + \mathbf{b}_A^{(M)} \right)
 \end{aligned} \tag{4.4}$$

where $\delta : \mathbb{R}^{d_m} \mapsto \mathbb{R}^{d_m}$ is a nonlinear activation function. Possible choices of δ include sigmoid, ReLU [47] and tanh functions.

The mechanism of AttNet is that, given the attribute vector $\mathbf{x}_k \in \mathbb{R}^u$ as the input, the first layer uses the weight matrix $\mathbf{W}_A^{(1)} \in \mathbb{R}^{d_1 \times u}$ and bias vector $\mathbf{b}_A^{(1)} \in \mathbb{R}^{d_1}$ to map \mathbf{x}_k to the output $\mathbf{V}_k^{(1)} \in \mathbb{R}^{d_1}$ with $d_1 < u$. The $\mathbf{V}_k^{(1)}$ is subsequently used as the input to the next layer. For simplicity, with

the d_M hidden units in the M -th layer, we denote the `AttNet` as

$$\Theta_A : \mathbb{R}^u \mapsto \mathbb{R}^{d_M} \quad (4.5)$$

Θ_A is parameterized by \mathbf{W}_A and \mathbf{b}_A , where $\mathbf{W}_A = (\mathbf{W}_A^{(1)}, \dots, \mathbf{W}_A^{(M)})$ and $\mathbf{b}_A = (\mathbf{b}_A^{(1)}, \dots, \mathbf{b}_A^{(M)})$.

The `SeqNet` is designed to learn the dependencies between items in the input sequences. `SeqNet` takes advantage of long short-term memory (LSTM) [25] network to learn both long and short-term dependencies within the sequences.

Given a sequence $\mathbf{S}_k \in \mathbb{R}^{T \times r}$ as the input, we have the parameters within `SeqNet` for each time step t as

$$\begin{aligned} \mathbf{i}_k^{(t)} &= \sigma \left(\mathbf{W}_i \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_i \mathbf{h}_k^{(t-1)} + \mathbf{b}_i \right) \\ \mathbf{f}_k^{(t)} &= \sigma \left(\mathbf{W}_f \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_f \mathbf{h}_k^{(t-1)} + \mathbf{b}_f \right) \\ \mathbf{o}_k^{(t)} &= \sigma \left(\mathbf{W}_o \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_o \mathbf{h}_k^{(t-1)} + \mathbf{b}_o \right) \\ \mathbf{g}_k^{(t)} &= \tanh \left(\mathbf{W}_c \boldsymbol{\alpha}_k^{(t)} + \mathbf{U}_c \mathbf{h}_k^{(t-1)} + \mathbf{b}_c \right) \\ \mathbf{c}_k^{(t)} &= \mathbf{f}_k^{(t)} \odot \mathbf{c}_k^{(t-1)} + \mathbf{i}_k^{(t)} \odot \mathbf{g}_k^{(t-1)} \\ \mathbf{h}_k^{(t)} &= \mathbf{o}_k^{(t)} \odot \tanh \left(\mathbf{c}_k^{(t)} \right) \end{aligned} \quad (4.6)$$

where $\sigma(z) = \frac{1}{1+e^{-z}}$ is a sigmoid activation function, \odot is the bitwise multiplication, $\mathbf{i}_k^{(t)}$, $\mathbf{f}_k^{(t)}$ and $\mathbf{o}_k^{(t)}$ are the internal gates of the LSTM, $\mathbf{c}_k^{(t)}$ and $\mathbf{h}_k^{(t)}$ are the cell and hidden states of the LSTM. For simplicity, we denote the SeqNet with d_S hidden units as

$$\Theta_S : \mathbb{R}^{T \times r} \mapsto \mathbb{R}^{d_S} \quad (4.7)$$

Θ_S is parameterized by bias vector set $\mathbf{b}_S = (\mathbf{b}_i, \mathbf{b}_f, \mathbf{b}_o, \mathbf{b}_c)$ and the set of weight matrices $\mathbf{W}_S = \{\mathbf{W}_{\theta_s}, \mathbf{U}_{\theta_s}\}$, where $\mathbf{W}_{\theta_s} = (\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_c) \in \mathbb{R}^{4 \times d_S \times r}$ and $\mathbf{U}_{\theta_s} = (\mathbf{U}_i, \mathbf{U}_f, \mathbf{U}_o, \mathbf{U}_c) \in \mathbb{R}^{4 \times d_S \times d_S}$.

4.2.2 FusionNet

Next, we design the FusionNet to integrate AttNet and SeqNet into one network. Here we propose three designs: (1) balanced, (2) AttNet-centric and (3) SeqNet-centric.

- **Balanced Design (Figure 4.1a)**. Both the attribute and sequence of each attributed sequence are processed simultaneously, and the results are concatenated together. The AttNet and SeqNet are first concatenated, followed by a fully connected layer with a nonlinear function over the concatenation to capture the dependencies between attributes and

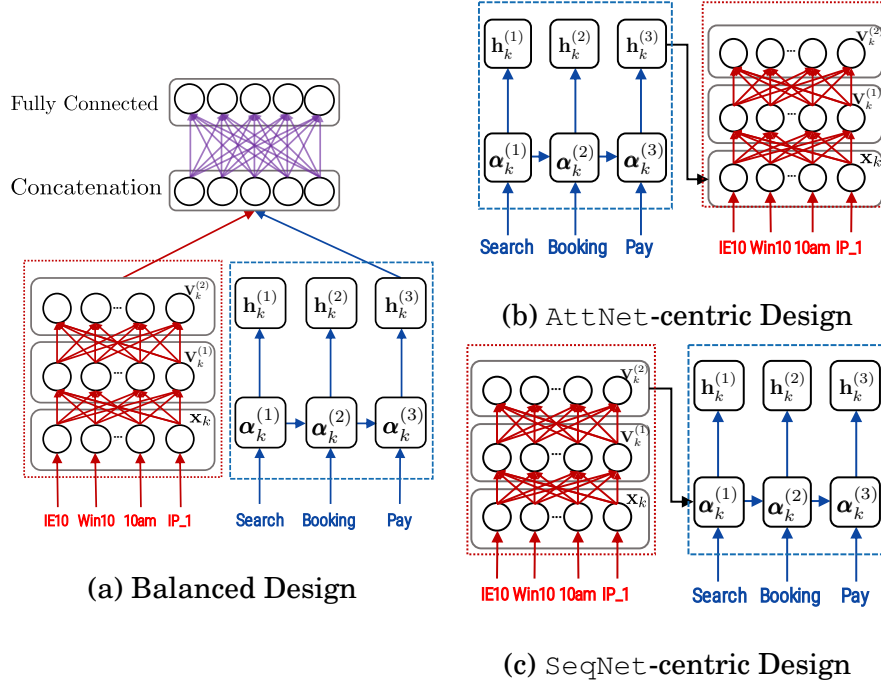


Figure 4.1: Comparison of three different network designs.

sequences. The output of this fully connected layer corresponds to the output of SeqNet after processing the last time step in the sequence input. We denote the balanced design as

$$\mathbf{y}_k = \mathbf{V}_k^{(M)} \oplus \mathbf{h}_k^{(T_k)} \quad (4.8)$$

$$\mathbf{z}_k = \delta(\mathbf{W}_z \mathbf{y}_k + \mathbf{b}_z) \quad (4.9)$$

where \oplus represents the concatenation operation, $\mathbf{W}_z \in \mathbb{R}^{d \times (d_M + d_S)}$ and $\mathbf{b}_z \in \mathbb{R}^d$ denote the weight matrix and bias vector in this fully connected layer, respectively.

• **AttNet-centric Design (Figure 4.1b)**. Here, the sequence is first transformed by sequence network, *i.e.*, the function Θ_S , and then used as an input of the attribute network, *i.e.*, the function Θ_A . Specifically, we modify Eq. 5.7 to incorporate sequence representation as an input. We use the output of SeqNet after processing the last time step in the sequence input. The modified $\mathbf{V}_k^{(1)}$ is written as

$$\mathbf{V}_k^{(1)} = \delta(\mathbf{W}_A^{(1)}(\mathbf{x}_k \oplus \mathbf{h}_k^{(T_k)}) + \mathbf{b}_A^{(1)}) \quad (4.10)$$

where the $\mathbf{W}_A^{(1)} \in \mathbb{R}^{d_1 \times (u+d_S)}$ and $\mathbf{b}_A^{(1)} \in \mathbb{R}^{d_1}$.

• **SeqNet-centric Design (Figure 4.1c)**. The attribute vector is first transformed by Θ_A and then used as an additional input for Θ_S . Specifically, we modify Eq. 6.3 to integrate attribute representations at the first time step as an input. The modified $\mathbf{h}_k^{(t)}$ is

$$\mathbf{h}_k^{(1)} = \mathbf{o}_k^{(1)} \odot \tanh(\mathbf{c}_k^{(1)}) + \mathbf{V}_k^{(M)} \quad (4.11)$$

In order to fuse AttNet and SeqNet using the SeqNet-centric design, the dimension of $\mathbf{V}_k^{(M)}$ has to be the same as $\mathbf{o}_k^{(1)}$ and $\mathbf{c}_k^{(1)}$. That is, $d_S = d_M$.

Without loss of generality, we summarize the above three designs as

$$\Theta : (\mathbb{R}^u, \mathbb{R}^{T \times r}) \mapsto \mathbb{R}^d \quad (4.12)$$

4.2.3 MetricNet

Without loss of generality, we present the `MetricNet` using the proposed balanced design due to space limitations. In the balanced design (as shown in Figure 4.1a), the explicit form of Eq. 4.12 can be written as

$$\Theta(J_k) = \Theta_A(\Theta_A(\mathbf{x}_k) \oplus \Theta_S(\mathbf{S}_k)) \quad (4.13)$$

Given an attributed sequence feedback instance (p_i, p_j, ℓ_{ij}) , where $p_i = (\mathbf{x}_i, \mathbf{S}_i)$ and $p_j = (\mathbf{x}_j, \mathbf{S}_j)$, $\ell_{ij} \in \{1, 0\}$. This input is transformed to $\Theta(p_i) \in \mathbb{R}^d$ and $\Theta(p_j) \in \mathbb{R}^d$ by the nonlinear transformation Θ .

The `MetricNet` is designed using a contrastive loss function [23] so that attributed sequences in each similar pair in \mathcal{S} have a smaller distance compared to those in \mathcal{D} after learning the distance metric. The `MetricNet` computes the Euclidean distance between each pair using the labels and back-propagates the gradients through all components in

our network. The learning objective of `MetricNet` can be written as

$$L(p_i, p_j, l_{ij}) = \frac{1}{2}(1 - l_{ij})(D_\Theta)^2 + \frac{1}{2}l_{ij}\{\max(0, g - D_\Theta)\}^2 \quad (4.14)$$

where g is the margin parameter, meaning that the pairs with a dissimilar label ($l_{ij} = 1$) contribute to the learning objective if and only if when the Euclidean distance between them is smaller than g [23]. We note that the `MetricNet` can augment all three designs in the same way. Figure 5.1 illustrates the `MetricNet` with the proposed balanced design.

4.3 Learning Feedback

In this section, we present the feedback learning mechanism of our MLAS network. Given two attributed sequences p_i and p_j as inputs, with Eq. 4.13, we have

$$\nabla L \equiv \left[\frac{\partial L}{\partial \mathbf{W}_A}, \frac{\partial L}{\partial \mathbf{b}_A}, \frac{\partial L}{\partial \mathbf{W}_S}, \frac{\partial L}{\partial \mathbf{b}_S} \right] \quad (4.15)$$

where the explicit form can be written as

$$\nabla L = \frac{\partial L}{\partial D_\Theta} \frac{\partial D_\Theta}{\partial \Theta} \left[\frac{\partial \mathbf{V}_k^{(M)}}{\partial \mathbf{W}_A}, \frac{\partial \mathbf{V}_k^{(M)}}{\partial \mathbf{b}_A}, \frac{\partial \mathbf{h}_k^{(T_k)}}{\partial \mathbf{W}_S}, \frac{\partial \mathbf{h}_k^{(T_k)}}{\partial \mathbf{b}_S} \right] \quad (4.16)$$

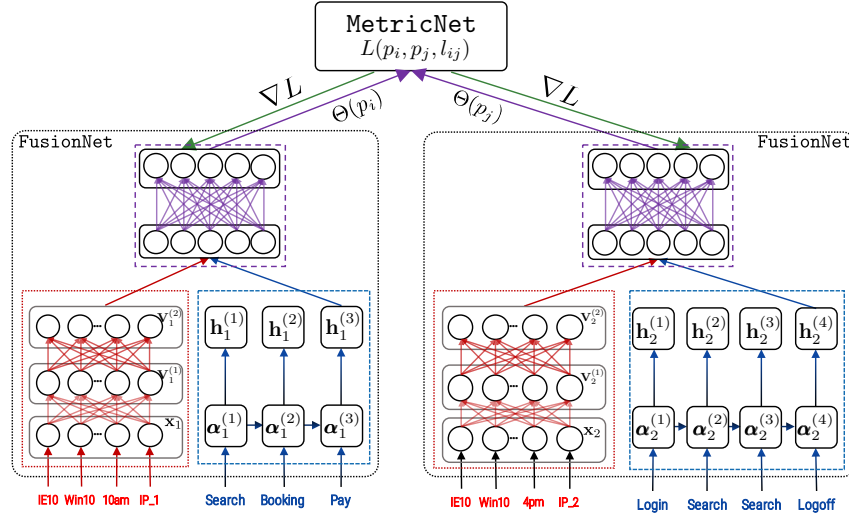


Figure 4.2: MLAS network with balanced design. Parameters in the two FusionNet are identical. Gradient ∇L is used to update all layers.

where

$$\frac{\partial L}{\partial D_{\Theta}} = (1 - l_{ij})D_{\Theta} - l_{ij} \max(0, g - D_{\Theta}) \quad (4.17)$$

$$\frac{\partial D_{\Theta}}{\partial \Theta} = \left(\Theta(p_i) - \Theta(p_j) \right) \cdot \left(\mathbf{1} - (\Theta(p_i) - \Theta(p_j)) \right) \quad (4.18)$$

where $\mathbf{1}$ is a vector filled with ones with the same shape as $\Theta(p_i)$ and $\Theta(p_j)$.

For the m -th layer in AttNet, we employ the following update functions

$$\begin{aligned} \frac{\partial \mathbf{V}_k^{(m)}}{\partial \mathbf{W}_A^{(m)}} &= \mathbf{V}_k^{(m)} \left(\mathbf{1} - \mathbf{V}_k^{(m)} \right) \mathbf{V}_k^{(m-1)} \\ \frac{\partial \mathbf{V}_k^{(m)}}{\partial \mathbf{b}_A^{(m)}} &= \mathbf{V}_k^{(m)} \left(\mathbf{1} - \mathbf{V}_k^{(m)} \right) \end{aligned} \quad (4.19)$$

With the learning rate γ , the parameters \mathbf{W}_A , \mathbf{W}_S , \mathbf{b}_A and \mathbf{b}_S can be updated by the following equations until convergence:

$$\begin{aligned}
 \mathbf{W}_A &= \mathbf{W}_A - \gamma \frac{\partial L}{\partial \mathbf{W}_A} \\
 \mathbf{b}_A &= \mathbf{b}_A - \gamma \frac{\partial L}{\partial \mathbf{b}_A} \\
 \mathbf{W}_S &= \mathbf{W}_S - \gamma \frac{\partial L}{\partial \mathbf{W}_S} \\
 \mathbf{b}_S &= \mathbf{b}_S - \gamma \frac{\partial L}{\partial \mathbf{b}_S}
 \end{aligned} \tag{4.20}$$

Algorithm 2 MLAS Learning

INPUT: A set of attributed sequences $\mathcal{J} = \{J_1, \dots, J_n\}$, a set of feedback as pairwise attributed sequences $\mathcal{C} = \{(p_i, p_j, \ell_{ij}) | p_i, p_j \in \mathcal{J}, \forall i, j = 1, \dots, n, i \neq j\}$, the number of layers M , learning rate γ , number of iterations ϑ and convergence error ϵ .

OUTPUT: Parameter sets $\{\mathbf{W}_A, \mathbf{b}_A, \mathbf{W}_S, \mathbf{b}_S\}$.

- 1: Initialize MLAS network Θ .
 - 2: **for each** $\vartheta' = 1, \dots, \vartheta$ **do**
 - 3: **for each** $(p_i, p_j, \ell_{ij}) \in \mathcal{C}$ **do**
 - 4: //Forward propagation.
 - 5: Calculate $\Theta(p_i)$ and $\Theta(p_j)$.
 - 6: Calculate D_Θ using Eq. 5.5.
 - 7: Calculate loss $L_{\vartheta'}(p_i, p_j, \ell_{ij})$ according to Eq. 4.14.
 - 8: **if** $|L_{\vartheta'}(p_i, p_j, \ell_{ij}) - L_{\vartheta'-1}(p_i, p_j, \ell_{ij})| < \epsilon$ **then**
 - 9: **break**
 - 10: **end if**
 - 11: //Back-propagation.
 - 12: Calculate $\frac{\partial L}{\partial \Theta}$ according to Eq. 4.17, 4.18.
 - 13: Calculate ∇L according to Eq. 5.14.
 - 14: Use Eq. 5.24 to update network parameters.
 - 15: **end for**
 - 16: **end for**
-

We summarize the algorithms for updating the MLAS network in Algorithm 3.

4.4 Experiments

4.4.1 Datasets

We evaluate the proposed methods using four real-world datasets. Two of them are derived from application log files¹ at Amadeus [2] (denoted as AMS-A and AMS-B). The other two datasets are derived from the Wikispeedia [68] dataset (denoted as Wiki-A and Wiki-B).

- **AMS-A:** We extracted $\sim 58k$ user sessions from log files of an internal application from our Amadeus. This internal application from Amadeus has motivated this research. Each record is composed of a user profile containing information ranging from system configurations to office name, and a sequence of functions invoked by click activities on the web interface. There are 288 distinct functions, 57,270 distinct user profiles in this dataset. The average length of the sequences is 18. 100 attributed sequence feedback pairs were selected by domain experts.

¹Personal identity information is not collected.

- **AMS-B:** There are $\sim 106\text{k}$ user sessions derived from internal application log files with 575 distinct functions and 106,671 distinct user profile. The average length of the sequences is 22. 84 attributed sequence feedback pairs were selected by domain experts.
- **Wiki-A:** This dataset is sampled from Wikispeedia dataset. Wikispeedia dataset originated from an online computation game [68], in which each user is given two pages (*i.e.*, source, and destination) from a subset of Wikipedia pages and asked to navigate from the source to the destination page. We use a subset of $\sim 2\text{k}$ paths from Wikispeedia with the average length of the path as 6. We also extract 11 sequence context as attributes (*e.g.*, the category of the source page, average time spent on each page, *etc.*). There are 200 feedback instances selected based on the criteria of frequent subsequences and attribute value.
- **Wiki-B:** This dataset is also sampled from Wikispeedia dataset. We use a subset of $\sim 1.5\text{k}$ paths from Wikispeedia with the average length of the path as 8. We also extract 11 sequence context (*e.g.*, the category of the source page, average time spent on each page,

Table 4.1: Summary of Compared Methods

Method	Data Used	Short Description	Reference
ATT	Attributes	Only <i>attribute</i> feedback is used in the model.	[26]
SEQ	Sequences	Only <i>sequence</i> feedback is used in the model.	[46]
ASF	Attributes Sequences	Feedback of attributes and sequences are used to train two models <i>separately</i> .	[26] + [46]
MLAS _{-B}	Attributes Sequences Dependencies	Balanced design using attributed sequence feedback to train one <i>unified</i> model.	This Work
MLAS _{-A}	Attributes Sequences Dependencies	Attribute-centric design using attributed sequence feedback to train one <i>unified</i> model.	This Work
MLAS _{-S}	Attributes Sequences Dependencies	Sequence-centric design using attributed sequence feedback to train one <i>unified</i> model.	This Work

etc) as attributes. 220 feedback instances have been selected based on the criteria of frequent subsequences and attribute value.

4.4.2 Compared Methods

We validate the effectiveness of our proposed MLAS solution compared with state-of-the-art baseline methods. To well understand the advancements of the proposed methods, we use baselines that are working on

only attributes (denoted as ATT) or sequences (denoted as SEQ), as well as methods without exploiting the dependencies between attributes and sequences (denoted as ASF). We summarize the compared methods and references in Table [4.1](#).

- Attribute-only Feedback (ATT) [\[26\]](#): Only attribute feedback is used in this model. This model first transforms fixed-size input data into feature vectors, then learns the similarities between these two inputs.
- Sequence-only Feedback (SEQ) [\[46\]](#): Only sequence feedback is used in this model. This model utilizes a long short-term memory (LSTM) to learn the similarities between two sequences.
- Attribute and Sequence Feedback (ASF) [\[26\]](#) + [\[46\]](#): This method is a combination of the ATT and SEQ methods as above, where the two networks are trained *separately* using attribute feedback and sequence feedback, respectively. The feature vectors generated by the two models are then concatenated.
- Balanced Network Design with Attributed Sequence Feedback ($MLAS-B$): The balanced design model using attributed sequence feedback to train a *unified* model.

- **AttNet-centric Network Design with Attributed Sequence Feedback (MLAS-A):** The AttNet-centric design using attributed sequence feedback to train a *unified* model.
- **SeqNet-centric Network Design with Attributed Sequence Feedback (MLAS-S):** The SeqNet-centric design using attributed sequence feedback to train a *unified* model.

4.4.3 Experimental Settings

4.4.3.1 Network Initialization and Training

Initializing the network parameters is important for models using gradient descent based approaches [16]. The weight matrices W_A in Θ_A and the W_S in Θ_S are initialized using the uniform distribution [20], the biases b_A and b_S are initialized with zero vector $\mathbf{0}$ and the recurrent matrix U_S is initialized using orthogonal matrix [55]. We use one hidden layer ($M = 1$) for AttNet and ATT in the experiments to make the training process easier.

After that, we pre-train each compared method. Pre-training is an important step to initialize the neural network-based models [16]. Our pre-training uses the attributed sequences as the inputs for FusionNet,

and use the generated feature representations to reconstruct the attributed sequence inputs. We also pre-train the `ATT` and `SEQ` networks in a similar fashion that reconstruct the respective attributes or sequences. We utilize ℓ_2 -regularization and early stopping strategy to avoid overfitting. Twenty percents of feedback pairs are used in the validation set. We choose `ReLU` activation function [47] in our `AttNet` to accelerate the stochastic gradient descent convergence.

4.4.3.2 Performance Evaluation Setting.

We evaluate the performance by using the feature representations generated by each method for clustering tasks. The feature representations are generated through a forward pass.

Clustering tasks have been widely used in distance metric learning work [70, 60]. In this set of experiments, we use `HDBSCAN` [8] clustering algorithm. `HDBSCAN` is a deterministic algorithm, which gives identical output when using the same input. We measure the normalized mutual information (NMI) [42] score. The maximum NMI score is 1. Specifically, we conduct the below two experiments:

1. The effect of feedback. We compare the performance of the clustering algorithm using the feature representations generated by FusionNet before and after incorporating the feedback.
2. The effect of varying parameters in the clustering algorithm. After the metric learning process, we evaluate the feature representations generated by all compared methods under various parameters of the clustering algorithm.

4.4.3.3 Parameter Study Settings.

We first evaluate the effect of output dimensions (*i.e.*, the dimension of the hidden layer), which not only affect the model size but also affects the performance of subsequent mining algorithms.

The other parameter we evaluate is the relative importance of attribute data (denoted as ω_A) in the attributed sequences. The pre-training phase is essential to gradient descent-based methods [16]. The relative importance of attribute data and sequence data are represented by the weights of Θ_A and Θ_S , denoted as ω_A and ω_S , where $\omega_A + \omega_S = 1$. The intuition is that with one data type more important, the other one becomes relatively less important.

4.4.4 Results and Analysis

4.4.4.1 Effect of Feedback.

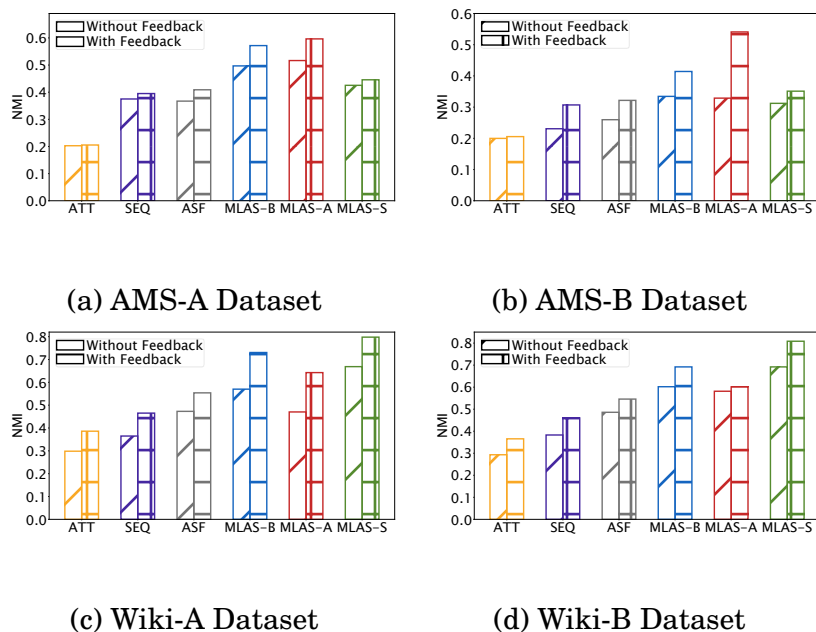


Figure 4.3: The effectiveness of using feedback. Using feedback could boost performance of all methods. The three methods we proposed (MLAS-B/A/S) are capable of exploiting the information of attributes, sequences, and more importantly, the attribute-sequence dependencies to outperform other methods.

We present the performance comparisons in clustering tasks using feature representations generated using the parameters of all methods in Figure 4.3. Two sets of feature representations are generated, the first set is generated after the pre-training (denoted as *without feed-*

back), the other set is generated after the metric learning step (denoted as *with feedback*). We fix the output dimension to 10, minimum cluster size to 100 and $\omega_A = 0.5$ (for MLAS-B/A/S). We have observed that the feedback can boost the performance of all methods, and the three methods (MLAS-B/A/S) proposed in this work are capable of outperforming other methods. Also, we also observe that the proposed three MLAS variations have better performance compared to the ASF, which also uses the information from attributes and sequences but *without* using the attribute-sequence dependencies.

Based on the above observations, we can conclude that the performance boost of our three architectures (MLAS-B/A/S) is a result of taking advantages of attribute data, sequence data, and more importantly, the attribute-sequence dependencies.

4.4.4.2 Performance in Clustering Tasks.

The primary parameter in HDBSCAN is the minimum cluster size [8], denoting the smallest set of instances to be considered as a group. Intuitively, while the minimum cluster size increases, each cluster may include instances that do not belong to it and the performance decreases.




Figure 4.4: Performance with varying clustering parameters. Clustering results using the feature representations produced by MLAS are the best among the compared methods.

Figure [4.4](#) presents the results with the output dimension is 10 and $\omega_A = 0.5$.

Compared to the best baseline method `ASF`, `MLAS-A` achieves up to 18.3% and 25.4% increase of performance on `AMS-A` and `AMS-B` datasets,

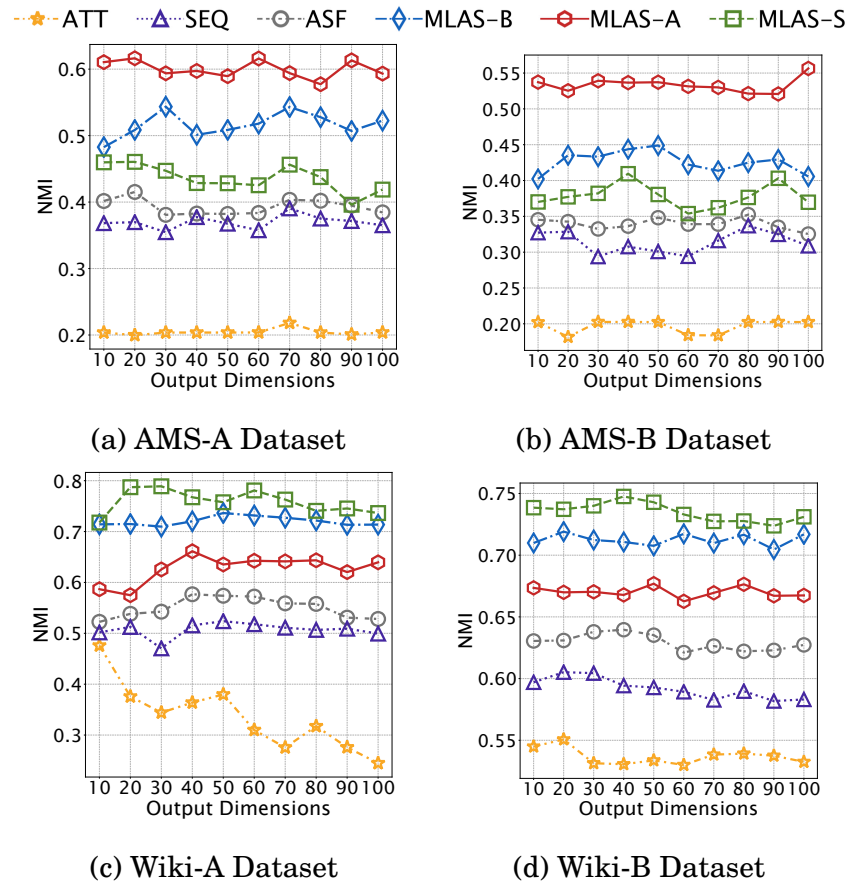


Figure 4.5: The effect of output dimensions (higher is better). Output dimension is an important factor for (1) the size of model; and (2) the cost of computations in downstream data mining tasks. Using the feature representations produced by MLAS can constantly achieve the best performance among the compared methods.

respectively. On Wiki-A and Wiki-B datasets, MLAS-S is capable of achieving up to 26.3% and 24.8% performance improvement compared to ASF, respectively. We can further confirm that MLAS network is capable

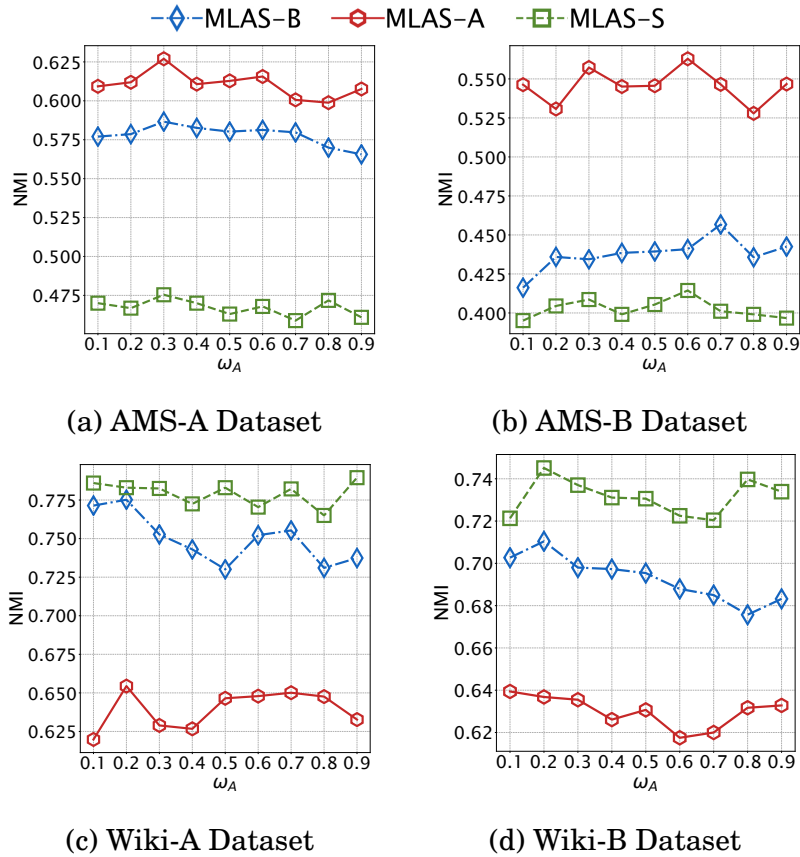


Figure 4.6: The effect of pre-training parameters in MLAS. The pre-training parameter ω_A decides the relative importance of attributes in the model. We observe that MLAS-A is capable of achieving the best performance on AMS-A and AMS-B datasets while MLAS-S has the best performance on Wiki-A and Wiki-B datasets.

of exploiting the *attribute-sequence dependencies* to improve the performance of the clustering algorithm with various parameter settings.

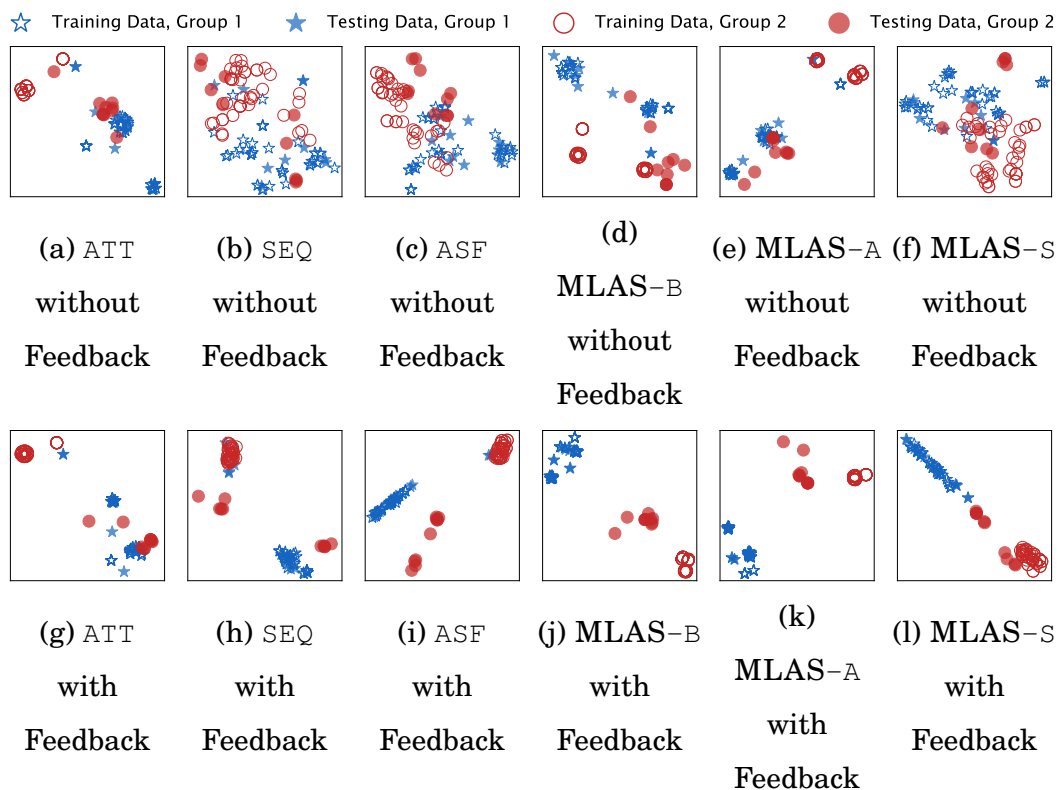


Figure 4.7: Plots of the feature representations. The MLAS is capable of exploiting the feedback and separating the instances from two different groups while keeping the instances from the same group together.

4.4.4.3 Output Dimensions

We evaluate MLAS under a wide range of output dimension choices. The number of output dimensions relates to a variety of impacts, such as the usability of feature representations in downstream data mining tasks. In this set of experiments, we fix the minimum cluster size at 50, $\omega_A = 0.5$ and vary output dimensions from 10 to 100. From Figure 4.5 we

conclude that our proposed approaches outperform the baseline methods with various output dimensions.

In particular, compared to the baseline method with the best performance, namely `ASF`, `MLAS-A` achieves 20.7% improvement on average on `AMS-A` dataset and 19.4% improvement on average on the `AMS-B` dataset. When evaluated using `Wiki-A` and `Wiki-B` datasets, `MLAS-S` outperforms `ASF` by 20.8% and 10.6% on average, respectively.

4.4.4.4 Pre-training Parameters

We evaluate `MLAS` under different pre-training parameters in this set of experiments. `ATT` and `SEQ` are not included in this set of experiments since they only utilize one data type. Output dimension is set to 5. Minimum cluster size is set at 50. Figure 4.6 presents the results under different pre-training parameters. This confirms that our proposed `MLAS` method is not sensitive to different pre-training parameters.

We notice the performance differences among the three `MLAS` architectures in the above experiments, where `MLAS-A` has the best performance on `AMS-A` and `AMS-B` datasets, and `MLAS-S` has the best performance on `Wiki-A` and `Wiki-B` datasets. We conclude this difference may relate to the datasets.

4.4.5 Case Studies

In Figure 4.7, we apply t-SNE [41] to the feature representations generated by all compared methods. The set of feature representations without feedback is generated after the pre-training phase and before the distance metric learning process.

Our goal is to demonstrate the differences in the feature space of each method. We randomly select data points from both training and testing sets with a ground truth of two groups. We have the following findings:

1. The methods using either attribute data (ATT) or sequence data (SEQ) *only* cannot use the attributed sequence feedback.
2. The method using *both* attributes and sequences *separately* (ASF) is capable of better separating the two groups than the methods using single data type (ATT and SEQ).
3. Our methods using attributed sequence feedback as a unity to train unified models (MLAS-B/A/S) are capable of separating the two groups the furthest, and thus achieve the best results.

These observations confirm that all three designs of MLAS can effectively learn the distance metric and result in better separation of two groups of data points.

Chapter 5

One-shot Learning on Attributed Sequences

5.1 Problem Definition

Inspired by the work in [6], we formulate our problem as finding the parameters θ of a predictor Θ that minimizes the loss $\mathcal{L}_{\text{one-shot}}$. Given a training set of g attributed sequences $\mathcal{G} = \{(p_1, c_1), \dots, (p_g, c_g)\}$, where each attributed sequence p_i has a unique class label c_i , we formulate the objective for one-shot learning for attributed sequences as:

$$\underset{\theta}{\text{minimize}} \sum_{(p_i, c_i) \in \mathcal{G}} \mathcal{L}_{\text{one-shot}}(\Theta(p_i; \theta), c_i) \quad (5.1)$$

That is, we want to minimize the loss calculated using the label predicted using parameter θ and the true label. One-shot learning is known as a hard problem [32] mainly as a result of unavoidable overfitting caused by insufficient data. With a complex data type, such as attributed sequences, the number of parameters that need to be trained is even larger, which further complicates the problem.

Table 5.1: Important Mathematical Notations

Notation	Description
\mathbb{R}	The set of real numbers
r	The number of possible items in sequences.
s_i	A sequence of categorical items.
$x_i^{(t)}$	The t -th item in sequence s_i .
t_{\max}	The maximum length of sequences in a dataset.
\mathbf{s}_i	A one-hot encoded sequence in the form of a matrix $\mathbf{s}_i \in \mathbb{R}^{t_{\max} \times r}$.
$\mathbf{x}_i^{(t)}$	A one-hot encoded item at t -th time step in a sequence.
\mathbf{v}_i	An attribute vector.
p_i	An attributed sequence. <i>i.e.</i> , $p_i = (\mathbf{v}_i, \mathbf{s}_i)$
\mathbf{p}_i	An n -dimensional feature vector of attributed sequence p_i .
Ω	A function transforming each attributed sequence to a feature vector.
d	A distance function. <i>e.g.</i> , Mahalanobis distance, Manhattan distance.
γ	An activation function within fully connected neural networks. Possible choices include <code>ReLU</code> and <code>tanh</code> .
σ	A logistic activation function within LSTM, <i>i.e.</i> , $\sigma(z) = \frac{1}{1+e^{-z}}$

5.2 The OLAS Model

5.2.1 Approach

In this work, we adopt an approach from the distance metric learning perspective. Distance metric learning methods are well known for several important applications, such as face recognition, image classification, *etc.* Distance metric learning is capable of disseminating data based on their dissimilarities using pairwise training samples. Recent work [32] has empirically demonstrated the effectiveness of the distance metric learning approach. In addition to the pairwise training samples,

there are two key components in distance metric learning: a *similarity* label depicting whether the training pair is similar and a distance function d . The similar and dissimilar pairs can be randomly generated using the class labels [32]. We define *attributed sequence triplets* in Definition 6.

Definition 6 (Attributed Sequence Triplets) An attributed sequence triplet (p_i, p_j, ℓ_{ij}) consists of two attributed sequences p_i, p_j , and a *similarity* label $\ell_{ij} \in \{0, 1\}$. The similarity label indicates whether p_i and p_j belong to the same class ($\ell_{ij} = 0$) or different classes ($\ell_{ij} = 1$). We denote $\mathcal{P} = \{(p_i, p_j, \ell_{ij}) | \ell_{ij} = 0\}$ as the *positive set* and $\mathcal{N} = \{(p_i, p_j, \ell_{ij}) | \ell_{ij} = 1\}$ as the *negative set*.

However, attributed sequences are not naturally represented as feature vectors. Therefore, we define a transformation function $\Omega(p_i; \omega)$ parameterized by ω as a part of the predictor Θ . Ω uses attributed sequences as the inputs and generates the corresponding feature vectors as the outputs. With two attributed sequences p_i and p_j as inputs, the

n -dimensional feature vectors of the respective attributed sequences are:

$$\begin{aligned}\mathbf{p}_i &= \Omega(p_i; \omega) \\ \mathbf{p}_j &= \Omega(p_j; \omega) \\ \mathbf{p}_i, \mathbf{p}_j &\in \mathbb{R}^n\end{aligned}\tag{5.2}$$

The other key component in distance metric learning approaches is a distance function (e.g., Mahalanobis distance [26], Manhattan distance [6]). A distance function is applied to the feature vectors in distance metric learning.

Distance metric learning-based approaches often use the Mahalanobis distance [26, 28], which can be equivalent to the Euclidean distance [26]. Using the two feature vectors of attributed sequences in Equation 5.2, the Mahalanobis distance can be written as:

$$d_\omega(\mathbf{p}_i, \mathbf{p}_j) = \sqrt{(\mathbf{p}_i - \mathbf{p}_j)^\top \Lambda (\mathbf{p}_i - \mathbf{p}_j)}\tag{5.3}$$

where d_ω is a specific form of distance function d denoting the inputs (i.e., $\mathbf{p}_i, \mathbf{p}_j$) are the results of transformations using parameter ω . $\Lambda \in \mathbb{R}^{n \times n}$ is a symmetric, semi-definite, and positive matrix, and Λ can be

decomposed as:

$$\Lambda = \Gamma^\top \Gamma, \quad (5.4)$$

where $\Gamma \in \mathbb{R}^{f \times n}$, $f \leq n$. By [70], Equation 5.3 is equivalent to:

$$\begin{aligned} d_\omega(\mathbf{p}_i, \mathbf{p}_j) &= \sqrt{(\mathbf{p}_i - \mathbf{p}_j)^\top \Gamma^\top \Gamma (\mathbf{p}_i - \mathbf{p}_j)} \\ &= \|\Gamma \mathbf{p}_i - \Gamma \mathbf{p}_j\|_2. \end{aligned} \quad (5.5)$$

Instead of directly minimizing the loss of the predictor function Θ predicting a label of each attributed sequence as in Equation 5.1, we can now achieve the same training goal by minimizing the loss of predicting whether a pair of attributed sequences belong to the same class using distance metric learning-based methods. The overall objective can be written as:

$$\underset{\omega}{\text{minimize}} \quad \sum_{(p_i, p_j, l_{ij}) \in \mathcal{P} \cup \mathcal{N}} \mathcal{L}(d_\omega(\mathbf{p}_i, \mathbf{p}_j), l_{ij}) \quad (5.6)$$

In recent work on distance metric learning applications [6, 26], deep neural networks are serve as the nonlinear transformation function Ω . Deep neural networks can effectively learn the features from input data without requiring domain-specific knowledge [32], and also generalize the knowledge for future predictions and inferences. These advantages make neural networks become an ideal solution for one-shot learning.

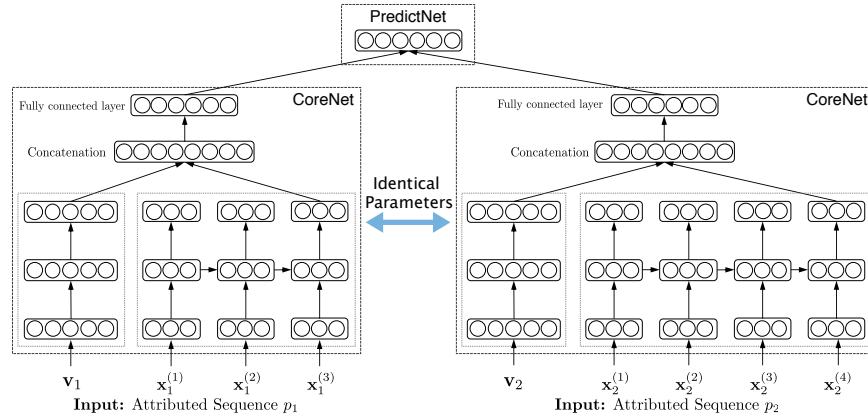


Figure 5.1: The network architecture of OLAS. The concatenation only happens after the *last time step* of the sequence so the information of the complete sequence is used.

5.2.2 OLAS Model Design

We next describe the design of the two key components of the OLAS model. First, we design a `CoreNet` for the nonlinear transformation of attributed sequences. Then, a `PredictNet` is designed to learn from the contrast of attributed sequences with different class labels. The specific parameters of the OLAS used in our experiments are detailed in Section [5.3](#).

The two main networks in `CoreNet`, a fully connected neural network with m layers and a long short-term memory (LSTM) network [\[25\]](#), correspond to the tasks of encoding the information from attributes and sequences in attributed sequences, respectively. By augmenting with another layer of fully connected neural network on top of the concate-

nation of the above networks, CoreNet is also capable of learning the attribute-sequence dependencies.

Given the input of an attribute vector $\mathbf{v}_k \in \mathbb{R}^u$, we define a fully connected neural network with m layers as:

$$\begin{aligned}
 \boldsymbol{\alpha}_1 &= \gamma(\mathbf{W}_1 \mathbf{v}_i + \mathbf{b}_1) \\
 \boldsymbol{\alpha}_2 &= \gamma(\mathbf{W}_2 \boldsymbol{\alpha}_1 + \mathbf{b}_2) \\
 &\vdots \\
 \boldsymbol{\alpha}_m &= \gamma(\mathbf{W}_m \boldsymbol{\alpha}_{m-1} + \mathbf{b}_m)
 \end{aligned} \tag{5.7}$$

where γ is a nonlinear transformation function. Although we use hyperbolic tangent \tanh in our model, other nonlinear functions such as rectified linear unit (ReLU) [47] can also be used depending on the empirical results. We denote the weights and bias parameters as:

$$\mathbf{W}_F = [\mathbf{W}_1, \dots, \mathbf{W}_m]^\top, \mathbf{b}_F = [\mathbf{b}_1, \dots, \mathbf{b}_m]^\top \tag{5.8}$$

Note that the choice of m is task-specific. Although neural networks with more layers are better at learning hierarchical structure in the data, it is also observed that such networks are challenging to train due to the

multiple nonlinear mappings that prevent the information and gradient passing along the computation graph [52].

W_F and b_F are used to transform the input of each layer to a lower dimension. This transformation is imperative given the often large number of dimensions of attribute vectors in real-world applications. Different from attribute vectors, the categorical items in the sequences in attributed sequences obey temporal ordering. The information of sequences is not only in the item values, but more importantly, in the temporal ordering of these items. In this vein, the CoreNet also utilizes an LSTM network. LSTM is capable of handling not only the ordering of items, but also the dependencies between different items in the sequences. Given a sequence s_i as the input, we use an LSTM [25] to

process each item $\mathbf{x}_k^{(t)}$ in this sequence as:

$$\begin{aligned}
\mathbf{i}^{(t)} &= \sigma \left(\mathbf{W}_i \mathbf{x}_k^{(t)} + \mathbf{U}_i \mathbf{h}^{(t-1)} + \mathbf{b}_i \right) \\
\mathbf{f}^{(t)} &= \sigma \left(\mathbf{W}_f \mathbf{x}_k^{(t)} + \mathbf{U}_f \mathbf{h}^{(t-1)} + \mathbf{b}_f \right) \\
\mathbf{o}^{(t)} &= \sigma \left(\mathbf{W}_o \mathbf{x}_k^{(t)} + \mathbf{U}_o \mathbf{h}^{(t-1)} + \mathbf{b}_o \right) \\
\mathbf{g}^{(t)} &= \tanh \left(\mathbf{W}_c \mathbf{x}_k^{(t)} + \mathbf{U}_c \mathbf{h}^{(t-1)} + \mathbf{b}_c \right) \\
\mathbf{c}^{(t)} &= \mathbf{f}^{(t)} \odot \mathbf{c}^{(t-1)} + \mathbf{i}^{(t)} \odot \mathbf{g}^{(t)} \\
\mathbf{h}^{(t)} &= \mathbf{o}^{(t)} \odot \tanh(\mathbf{c}^{(t)})
\end{aligned} \tag{5.9}$$

where σ is a sigmoid activation function, \odot denotes the bitwise multiplication, $\mathbf{i}^{(t)}$, $\mathbf{f}^{(t)}$ and $\mathbf{o}^{(t)}$ are the internal gates of the LSTM, $\mathbf{c}^{(t)}$ and $\mathbf{h}^{(t)}$ are the cell and hidden states of the LSTM. Without loss of generality, we denote LSTM kernel parameters \mathbf{W}_L , recurrent parameters \mathbf{U}_L and bias parameters \mathbf{b}_L as:

$$\begin{aligned}
\mathbf{W}_L &= [\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_c]^\top \\
\mathbf{U}_L &= [\mathbf{U}_i, \mathbf{U}_f, \mathbf{U}_o, \mathbf{U}_c]^\top \\
\mathbf{b}_b &= [\mathbf{b}_i, \mathbf{b}_f, \mathbf{b}_o, \mathbf{b}_c]^\top
\end{aligned} \tag{5.10}$$

The attribute vectors and sequences are processed simultaneously and the outputs of both networks are concatenated together. Instead of using

the outputs of the LSTM at every time step, we only concatenate the last output from the LSTM to the output of the fully connected neural network so that the complete sequence information is used. After that, another layer of the fully connected neural network is used to capture the dependencies between attributes and sequences. Given the output dimensions of α_m and $\mathbf{h}^{(t)}$ as n_m and n_l , respectively, the concatenation and the last fully connected layer of CoreNet can be written as:

$$\mathbf{p}_i = \gamma (\mathbf{W}_p (\alpha_m \oplus \mathbf{h}^{(t_i)}) + \mathbf{b}_p) \quad (5.11)$$

where \oplus represents the concatenation of two vectors, $\mathbf{W}_p \in \mathbb{R}^{n \times (n_m + n_l)}$ and $\mathbf{b}_p \in \mathbb{R}^n$ denote the weight matrix and bias vector in this fully connected layer for an n -dimensional output. In summary, the CoreNet can be written as:

$$\Omega : (\mathbb{R}^u, \mathbb{R}^{t_{\max} \times r}) \mapsto \mathbb{R}^n \quad (5.12)$$

The two outputs of CoreNet (\mathbf{p}_i and \mathbf{p}_j) are first generated. Then, \mathbf{p}_i , \mathbf{p}_j and the similarity label ℓ_{ij} , are used by the PredictNet to learn the similarities and differences between them. The PredictNet is designed to utilize a contrastive loss function [23] so that attributed sequences in different categories are disseminated. The contrastive loss function is composed of two parts: a partial loss for the dissimilar pairs and a

Algorithm 3 Training using attributed sequence triplets

INPUT: A positive set \mathcal{P} and a negative set \mathcal{N} of attributed sequence triplets, the number of layers in fully connected neural networks m , learning rate λ , number of iterations ϕ and convergence error ϵ .

OUTPUT: Parameters of OLAS ($\{\mathbf{W}_F, \mathbf{b}_F, \mathbf{W}_L, \mathbf{U}_L, \mathbf{b}_L\}$).

```

1: Initialize OLAS network.
2: for each  $\phi' = 1, \dots, \phi$  do           ▷  $\phi$  is the maximum number of training
   epochs.
3:   for each  $(p_i, p_j, l_{ij}) \in \mathcal{P} \cup \mathcal{N}$  do
4:      $\mathbf{p}_i \leftarrow \Omega(p_i; \omega)$ .
5:      $\mathbf{p}_j \leftarrow \Omega(p_j; \omega)$ .
6:     Compute  $d_\omega$ .                               ▷ Equation 5.5.
7:     Compute the loss  $\mathcal{L}_{\phi'}(\mathbf{p}_i, \mathbf{p}_j, l_{ij})$ .       ▷ Equation 5.13.
8:     if  $|\mathcal{L}_{\phi'}(\mathbf{p}_i, \mathbf{p}_j, l_{ij}) - \mathcal{L}_{\phi'-1}(\mathbf{p}_i, \mathbf{p}_j, l_{ij})| < \epsilon$  then
9:       break                                       ▷ Early stopping to avoid overfitting.
10:    else
11:      Compute  $\frac{\partial \mathcal{L}}{\partial d_\omega}, \frac{\partial d_\omega}{\partial \Omega}$ .           ▷ Equation 5.16, 5.17.
12:      Compute  $\nabla \mathcal{L}$ .                               ▷ Equation 5.14.
13:      Update network parameters.                   ▷ Equation 5.24.
14:    end if
15:  end for
16: end for

```

partial loss for similar pairs. The specific form of contrastive loss of PredictNet can be written as:

$$\begin{aligned}
 \mathcal{L}(\mathbf{p}_i, \mathbf{p}_j, l_{ij}) = & \underbrace{\frac{1}{2} l_{ij} \left[\max(0, \xi - d_\omega(\mathbf{p}_i, \mathbf{p}_j)) \right]^2}_{\text{Partial loss for dissimilar pairs.}} \\
 & + \underbrace{\frac{1}{2} (1 - l_{ij}) d_\omega^2(\mathbf{p}_i, \mathbf{p}_j)}_{\text{Partial loss for similar pairs}}
 \end{aligned} \tag{5.13}$$

where ξ is a margin parameter used to prevent the dataset being reduced to a single point [70]. That is, the attributed sequences with $\ell_{ij} = 1$ are only used to adjust the parameters in the transformation function Ω if the distance between them is larger than ξ . The architecture of OLAS is illustrated in Figure 5.1.

5.2.3 OLAS Model Training

With the contrastive loss \mathcal{L} computed using Equation 5.13, we can now calculate the gradient $\nabla\mathcal{L}$, which is used to adjust parameters in the network as:

$$\nabla\mathcal{L} \equiv \left[\frac{\partial\mathcal{L}}{\partial\mathbf{W}_F}, \frac{\partial\mathcal{L}}{\partial\mathbf{b}_F}, \frac{\partial\mathcal{L}}{\partial\mathbf{W}_L}, \frac{\partial\mathcal{L}}{\partial\mathbf{U}_L}, \frac{\partial\mathcal{L}}{\partial\mathbf{b}_L} \right] \quad (5.14)$$

With the transformation function Ω and distance function d , the explicit form of $\nabla\mathcal{L}$ can be written as:

$$\nabla\mathcal{L} = \frac{\partial\mathcal{L}}{\partial d_\omega} \frac{\partial d_\omega}{\partial\Omega} \left[\frac{\partial\boldsymbol{\alpha}_m}{\partial\mathbf{W}_F}, \frac{\partial\boldsymbol{\alpha}_m}{\partial\mathbf{b}_F}, \frac{\partial\mathbf{h}^{(t_i)}}{\partial\mathbf{W}_L}, \frac{\partial\mathbf{h}^{(t_i)}}{\partial\mathbf{U}_L}, \frac{\partial\mathbf{h}^{(t_i)}}{\partial\mathbf{b}_L} \right] \quad (5.15)$$

where

$$\begin{aligned} \frac{\partial\mathcal{L}}{\partial d_\omega} = & -\ell_{ij} \max(0, \xi - d_\omega(\mathbf{p}_i, \mathbf{p}_j)) \\ & + (1 - \ell_{ij}) d_\omega(\mathbf{p}_i, \mathbf{p}_j) \end{aligned} \quad (5.16)$$

$$\frac{\partial d_\omega}{\partial \Omega} = (\mathbf{p}_i - \mathbf{p}_j) \cdot (\mathbf{1} - (\mathbf{p}_i - \mathbf{p}_j)) \quad (5.17)$$

where $\mathbf{1}$ is a vector filled with ones.

For the m -th layer in a fully connected neural network, we employ the following update functions:

$$\begin{aligned} \frac{\partial \boldsymbol{\alpha}_m}{\partial \mathbf{W}_m} &= \boldsymbol{\alpha}_m (\mathbf{1} - \boldsymbol{\alpha}_m) \boldsymbol{\alpha}_{m-1} \\ \frac{\partial \boldsymbol{\alpha}_m}{\partial \mathbf{b}_m} &= \boldsymbol{\alpha}_m (\mathbf{1} - \boldsymbol{\alpha}_{m-1}) \end{aligned} \quad (5.18)$$

Here we use three steps to explain how OLAS back-propagates the gradients. We use a $\delta_{\mu,\nu}$ function to simplify the equations with $\mu = \{\mathbf{i}, \mathbf{f}, \mathbf{o}\}$ and $\nu = \{\mathbf{i}, \mathbf{f}, \mathbf{o}, \mathbf{c}\}$:

$$\delta_{\mu,\nu} = \begin{cases} 1, & \text{if } \mu = \nu \\ 0, & \text{otherwise} \end{cases} \quad (5.19)$$

First, we have the following equations for $\mathbf{h}^{(t)}$ and $\mathbf{c}^{(t)}$:

$$\begin{aligned}\frac{\partial \mathbf{h}^{(t)}}{\partial \mathbf{W}_\nu} &= \frac{\partial \mathbf{o}^{(t)}}{\partial \mathbf{W}_\nu} \odot \tanh(\mathbf{c}^{(t)}) + \mathbf{o}^{(t)} \\ &\quad \odot (1 - \tanh^2(\mathbf{c}^{(t)})) \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{W}_\nu} \\ \frac{\partial \mathbf{h}^{(t)}}{\partial \mathbf{U}_\nu} &= \frac{\partial \mathbf{o}^{(t)}}{\partial \mathbf{U}_\nu} \odot \tanh(\mathbf{c}^{(t)}) + \mathbf{o}^{(t)} \\ &\quad \odot (1 - \tanh^2(\mathbf{c}^{(t)})) \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{U}_\nu}\end{aligned}\tag{5.20}$$

$$\begin{aligned}\frac{\partial \mathbf{h}^{(t)}}{\partial \mathbf{b}_\nu} &= \frac{\partial \mathbf{o}^{(t)}}{\partial \mathbf{b}_\nu} \odot \tanh(\mathbf{c}^{(t)}) + \mathbf{o}^{(t)} \\ &\quad \odot (1 - \tanh^2(\mathbf{c}^{(t)})) \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{b}_\nu} \\ \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{W}_\nu} &= \frac{\partial \mathbf{f}^{(t)}}{\partial \mathbf{W}_\nu} \odot \mathbf{c}^{(t-1)} + \mathbf{f}^{(t)} \odot \frac{\partial \mathbf{c}^{(t-1)}}{\partial \mathbf{W}_\nu} + \\ &\quad \frac{\partial \mathbf{i}^{(t)}}{\partial \mathbf{W}_\nu} \odot \mathbf{g}^{(t)} + \mathbf{i}^{(t)} \odot \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{W}_\nu} \\ \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{U}_\nu} &= \frac{\partial \mathbf{f}^{(t)}}{\partial \mathbf{U}_\nu} \odot \mathbf{c}^{(t-1)} + \mathbf{f}^{(t)} \odot \frac{\partial \mathbf{c}^{(t-1)}}{\partial \mathbf{U}_\nu} + \\ &\quad \frac{\partial \mathbf{i}^{(t)}}{\partial \mathbf{U}_\nu} \odot \mathbf{g}^{(t)} + \mathbf{i}^{(t)} \odot \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{U}_\nu} \\ \frac{\partial \mathbf{c}^{(t)}}{\partial \mathbf{b}_\nu} &= \frac{\partial \mathbf{f}^{(t)}}{\partial \mathbf{b}_\nu} \odot \mathbf{c}^{(t-1)} + \mathbf{f}^{(t)} \odot \frac{\partial \mathbf{c}^{(t-1)}}{\partial \mathbf{b}_\nu} + \\ &\quad \frac{\partial \mathbf{i}^{(t)}}{\partial \mathbf{b}_\nu} \odot \mathbf{g}^{(t)} + \mathbf{i}^{(t)} \odot \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{b}_\nu}\end{aligned}\tag{5.21}$$

Then, we have the following equations for $\mathbf{i}^{(t)}$, $\mathbf{f}^{(t)}$ and $\mathbf{o}^{(t)}$:

$$\begin{aligned}
 \frac{\partial \Delta_\mu}{\partial \mathbf{W}_\nu} &= \Delta_\mu (1 - \Delta_\mu) \boldsymbol{\alpha}^{(t)} \delta_{\mu,\nu} \\
 \frac{\partial \Delta_\mu}{\partial \mathbf{U}_\nu} &= \Delta_\mu (1 - \Delta_\mu) \mathbf{h}^{(t-1)} \delta_{\mu,\nu} \\
 \frac{\partial \Delta_\mu}{\partial \mathbf{b}_\nu} &= \Delta_\mu (1 - \Delta_\mu) \delta_{\mu,\nu}
 \end{aligned} \tag{5.22}$$

where $\Delta_{\mathbf{i}} = \mathbf{i}^{(t)}$, $\Delta_{\mathbf{f}} = \mathbf{f}^{(t)}$ and $\Delta_{\mathbf{o}} = \mathbf{o}^{(t)}$.

Finally, we have the gradients for $\mathbf{g}^{(t)}$ as:

$$\begin{aligned}
 \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{W}_\nu} &= (1 - (\mathbf{g}^{(t)})^2) \boldsymbol{\alpha}^{(t)} \delta_{c,\nu} \\
 \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{U}_\nu} &= (1 - (\mathbf{g}^{(t)})^2) \mathbf{h}^{(t-1)} \delta_{c,\nu} \\
 \frac{\partial \mathbf{g}^{(t)}}{\partial \mathbf{b}_\nu} &= (1 - (\mathbf{g}^{(t)})^2) \delta_{c,\nu}
 \end{aligned} \tag{5.23}$$

With the learning rate λ , the parameters \mathbf{W}_F , \mathbf{W}_L , \mathbf{U}_L , \mathbf{b}_F and \mathbf{b}_L can be updated by the following equation until convergence is achieved:

$$\begin{aligned}
 \mathbf{W}_F &= \mathbf{W}_F - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{W}_F} \\
 \mathbf{b}_F &= \mathbf{b}_F - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{b}_F} \\
 \mathbf{W}_L &= \mathbf{W}_L - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{W}_L} \\
 \mathbf{U}_L &= \mathbf{U}_L - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{U}_L} \\
 \mathbf{b}_L &= \mathbf{b}_L - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{b}_L}
 \end{aligned} \tag{5.24}$$

We summarize the algorithms for updating the OLAS network in Algorithm 3.

5.2.4 Labeling Attributed Sequences

Once we have trained the OLAS network to recognize the similarities and dissimilarities between exemplars of attributed sequence pairs. The OLAS is then ready to be used to assign labels to unlabeled attributed sequences in one-shot learning. Given a test attributed sequence p_k from a set \mathcal{K} of unlabeled instances, a set $\mathcal{G} = \{p_g\}_{g=1}^G$ of attributed sequences with G categories, in which there is only one instance per category, and

the goal is to classify p_k into one of G categories. We can now use the OLAS network with only one forward pass to calculate the distance between p_k with each of the G attributed sequences and the label of the instance that is closest to p_k is then assigned as the label of p_k . This process can be defined using maximum similarity as:

$$\hat{c}_k = \arg \min_g d_\omega(\mathbf{p}_k, \mathbf{p}_g) \quad (5.25)$$

where \hat{c}_k is the predicted label of \mathbf{p}_k .

5.3 Experiments

5.3.1 Datasets

Our solution has been motivated in part by use case scenarios observed at Amadeus related to attributed sequences. For this reason, we now work with the log files of an Amadeus [2] internal application. Also, we apply our methodology to real-world, public available Wikispeedia data [68]. We summarize the data descriptions as follows:

- **Amadeus data (AMS1~AMS6).** We sampled six datasets from the log files of an internal application at Amadeus IT Group. Each

attributed sequence is composed of a user profile containing information (*e.g.*, system configuration, office name) and a sequence of function names invoked by web click activities (*e.g.*, login, search) ordered by time.

- **Wikispeedia data (WS1~WS6)**. Wikispeedia is an online game requiring participants to click through from a given start page to an end page using fewest clicks [68]. We select the *finished* path and extract several properties of each path (*e.g.*, the category of the start path, time spent per click). We also sample six datasets from Wikispeedia. The Wikispeedia data is available through the Stanford Network Analysis Project¹ [37].

Following the protocols in recent work [32], we utilize the attributed sequences associated with 60% of categories to generate attributed sequence triplets and use them in training.

The class labels used in training and one-shot learning are disjoint sets. Similar to the strategy in [34], where the authors designed a 20-way classification task that attempts to match an alphabet with one of the twenty possible classes, we randomly select one instance in the one-shot learning set and attempt to give it a correct label. We selected

¹<https://snap.stanford.edu/data/wikispeedia.html>

Table 5.2: Number of Classes in Datasets

Dataset	Training	One-shot Learning
AMS1, WS1	6	4
AMS2, WS2	12	8
AMS3, WS3	18	12
AMS4, WS4	24	16
AMS5, WS5	30	20
AMS6, WS6	36	24

Table 5.3: Compared Methods

Name	Data Used	Note
OLAS	Attributed Sequences	This Work
OLASEmb	Attributed Sequence Embeddings	This work + [79]
ATT	Attributes Only	[32]
SEQ	Sequence Only	[57] + [32]

2000 instances for each set used in one-shot learning and compute the accuracy. We summarize the number of classes in Table 5.2.

5.3.2 Compared Methods

We focus on one-shot learning methods on different data types. We summarize the compared methods in Table 6.2. Specifically, we compare the performance of the following one-shot learning methods:

- OLAS: We first evaluate our proposed method using attributed sequences data.
- OLASEmb: Instead of using the attributed sequence instances as input, we use the embeddings of attributed sequences as the input.

Algorithm 4 One-shot learning for attributed sequences.

INPUT: Trained networks `CoreNet` Ω and `PredictNet`, a set of unlabeled attributed sequences \mathcal{K} , a set of labeled attributed sequence with one example per class \mathcal{G} and a distance function d .

OUTPUT: A set of labeled attributed sequences \mathcal{K}' .

```

1:  $\mathcal{K}' \leftarrow \emptyset$ 
2: for each  $p_k \in \mathcal{K}$  do
3:    $\varepsilon \leftarrow +\infty$  ▷ Set initial minimum distance to  $+\infty$ 
4:    $\mathbf{p}_k \leftarrow \Omega(p_k; \omega)$ 
5:   for each  $(p_g, c_g) \in \mathcal{G}$  do
6:      $\mathbf{p}_g \leftarrow \Omega(p_g; \omega)$ 
7:     if  $d(\mathbf{p}_k, \mathbf{p}_g) \leq \varepsilon$  then
8:        $\varepsilon \leftarrow d_\omega(\mathbf{p}_k, \mathbf{p}_g)$  ▷ Using PredictNet.
9:        $\hat{c}_k \leftarrow c_g$  ▷ Assign the same label of  $p_g$  to  $p_k$ .
10:    end if
11:     $\mathcal{K}' \leftarrow (p_k, \hat{c}_k)$ 
12:  end for
13: end for
14: return  $\mathcal{K}$ 

```

We want to find out whether a simpler heuristic combination of state-of-the-art would achieve better performance.

- ATT: This is the state-of-the-art method [32] using only attributes of the data.
- SEQ: We combine the state-of-the-art in one-shot learning [32] with sequence-to-sequence learning [57] to be able to utilize sequences in one-shot learning.

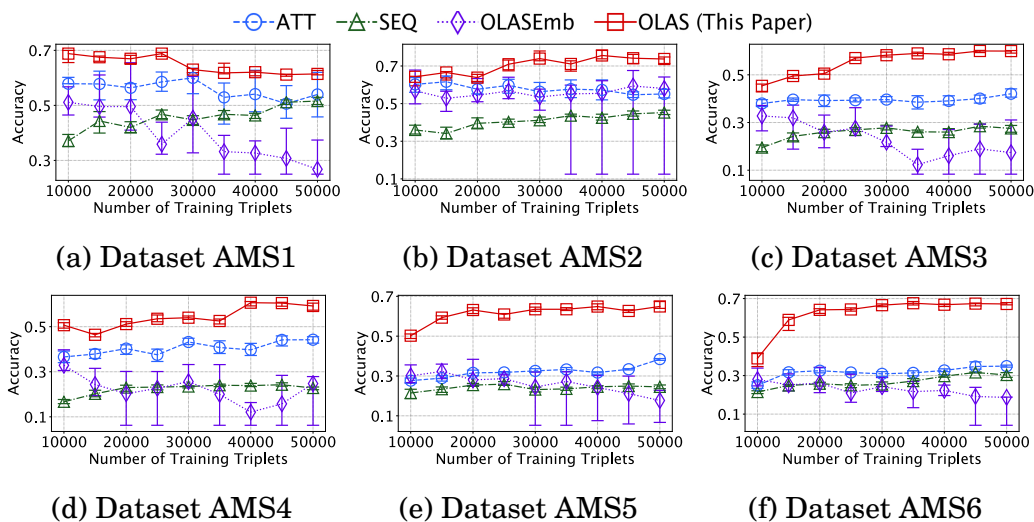


Figure 5.2: Accuracy of the label prediction on AMS datasets using **Euclidean** distance function.

5.3.3 Experiment Settings

5.3.3.1 Protocols

The goal of one-shot learning is to correctly assign class labels to each instance. In order to compare with state-of-the-art work [32, 6], we also use accuracy to evaluate the performance. A higher accuracy score means a method could make more correct class label predictions. For each experiment setting, we repeat ten times and report the median, 25 percentile and 75 percentile of the results using error bars. For each training process using attributed sequence triplets, we hold out 20% of the training data as the validation set. The holdout portion is not lim-

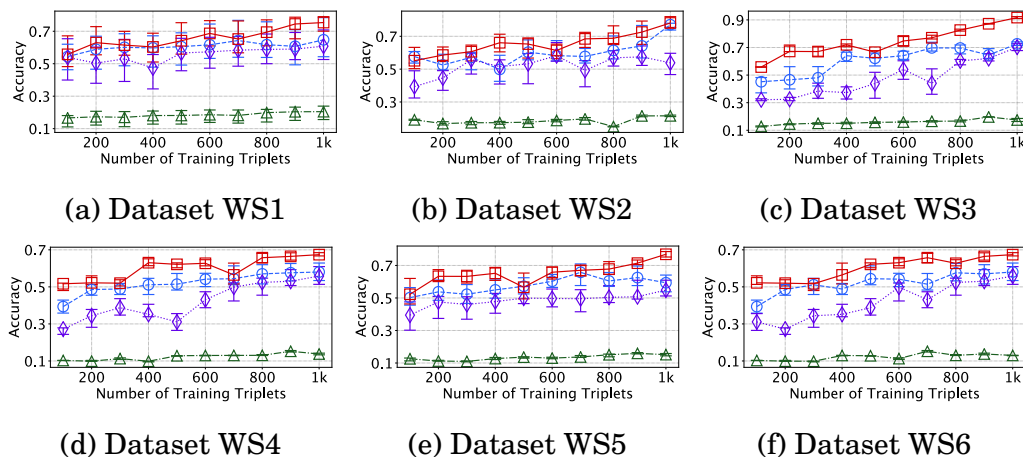


Figure 5.3: Accuracy of the label prediction on Wikispeedia datasets using **Euclidean** distance function.

ited to the instances with certain labels, but instead, they are randomly chosen from all possible classes.

5.3.3.2 Network Initialization and Settings

Gradient-based methods often require a careful initialization of the neural networks. In our experiments, we use normalized random distribution [20] to initialize weight matrices W_F and W_L , orthogonal matrix is used to initialize recurrent matrices U_L and biases are initialized to zero vector $\mathbf{0}$. Specifically, the m -th layer of the fully connected neural network is initialized as:

$$W_m \sim \text{Uniform} \left[-\frac{\sqrt{6}}{\sqrt{n_{m-1} + n_{m+1}}}, \frac{\sqrt{6}}{\sqrt{n_{m-1} + n_{m+1}}} \right]$$

where n_m is the output dimension of the m -th layer. There are three layers used in our experiments. Meanwhile, the weight matrices $\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_c$ are initialized as:

$$\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_c \sim \text{Uniform} \left[-\sqrt{\frac{6}{n_l}}, \sqrt{\frac{6}{n_l}} \right]$$

where n_l is the dimension of the output. In our experiments, we use 50 dimensions for both n_l and n_m . We utilize ℓ_2 -regularization with early stopping to avoid overfitting. The validation set is composed of 20% of the total amount of attributed sequence triplets in the training set.

5.3.3.3 Performance Studies

In this section, we present the performance studies of the proposed OLAS network and compare it with techniques in the state-of-the-art.

5.3.3.4 Varying number of training triplets

Figure [5.2](#) and [5.3](#) present the results where each setting has a fixed number of labels while the number of training triplets increases. Based on the experiment result figures, we have the following observations:

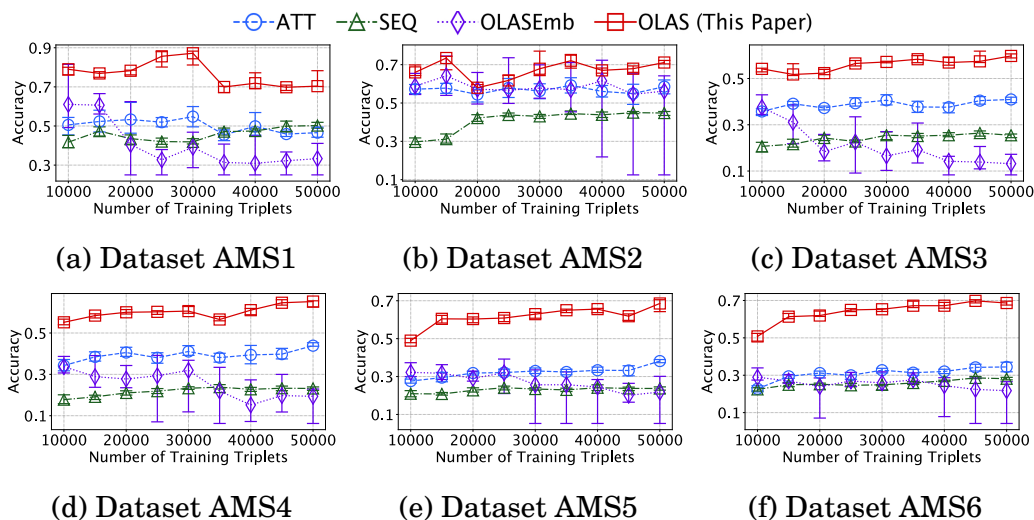


Figure 5.4: Accuracy of the label prediction on AMS datasets using **Manhattan** distance function.

- As more triplets being used in the training process, the accuracy of one-shot learning keeps increasing with the trained OLAS network. Intuitively, with more examples demonstrated to the OLAS, it could better gain a better capability of generalization, even though the data instances used in one-shot learning are previously unseen.
- Overfitting challenges the performance of all one-shot learning approaches. Although we use early stopping and ℓ_2 -regularization in all experiments, overfitting can still be challenging due to there is only one example per class in one-shot learning.
- OLAS can achieve better performance than other baseline methods when there are more possible classes. While OLAS maintains

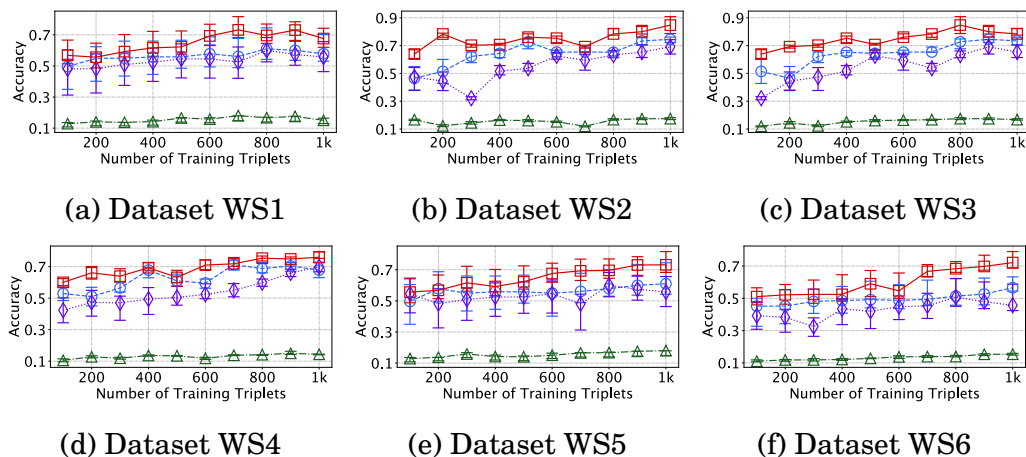


Figure 5.5: Accuracy of the label prediction on Wikispeedia datasets using **Manhattan** distance function.

a stable performance outperforming state-of-the-art under various parameter settings, OLAS can achieve a better performance when the classification task become *harder* with more possible class labels.

5.3.3.5 Observations using different datasets

Different from the synthetic datasets, the real-world applications often consist of diverse and noisy data instances. It is also interesting to examine the results using different real-world datasets. We find that the performance of OLAS remains superior when we use the twelve datasets sampled from two real-world applications.

5.3.3.6 Advantage of the *end-to-end* model

Although it is possible to use attributed sequence embeddings [79] with one-shot learning, the experiment results have proven that the performance of the end-to-end solution in this work is far superior to and more stable than OLASEmb. Specifically, the performance of our closest baseline method OLASEmb has varied more compared to all other methods. Building an end-to-end model allows the back-propagation of gradient throughout all layers in the OLAS model. On the other hand, the two gradients in OLASEmb, *i.e.*, the gradient in the model for generating attributed sequence embedding and the gradient in one-shot learning model, are independent and thus the parameters within this method cannot be better adjusted than our solution OLAS.

5.3.3.7 Effect of different distance functions.

Recent work [6] has observed significant differences in performance when using different distance functions. Here, we substitute the Euclidean distance function with the Manhattan distance function to see the performance of all compared methods. We observe that the proposed OLAS model is capable of achieving the best results despite which one of the two distance functions are used.

Chapter 6

Attention Model for Attributed Sequence Classification

6.1 Problem Definition.

We formulate the attributed sequence classification problem as the problem of finding the parameters θ of a predictor Θ that minimizes the prediction error of class labels. Intuitively, we want to maximize the possibility of correctly predicting labels when given a training set $\mathbb{P} = \{p_1, \dots, p_k\}$ of k attributed sequences. Thus, we formulate the training process as an optimization process:

$$\arg \min_{\theta} - \sum_i \Pr(l_i) \log \Pr(\Theta(p_i)) \quad (6.1)$$

Our goal is to find the parameters that minimize the categorical cross-entropy loss between the predicted labels using parameters in function Θ and the true labels for all attributed sequences in the dataset.

6.2 Attributed Sequence Attention Mechanism

The proposed AMAS model has three components, one `AttNet` for learning the attribute information, one `SeqNet` to learn the sequential infor-

Table 6.1: Important Mathematical Notations

Notation	Description
\mathbb{R}	The set of real numbers.
\mathbb{P}	A set of attributed sequences.
r	The number of all possible items in sequences.
s_i	A sequence of categorical items.
$x_i^{(t)}$	The t -th item in sequence s_i .
t_{\max}	The maximum length of sequences.
\mathbf{s}_i	A one-hot encoded sequence in the form of a matrix $\mathbf{s}_i \in \mathbb{R}^{t_{\max} \times r}$.
$\mathbf{x}_i^{(t)}$	A one-hot encoded item at t -th time step.
\mathbf{v}_i	An attribute vector.
p_i	An attributed sequence. <i>i.e.</i> , $p_i = (\mathbf{v}_i, \mathbf{s}_i)$
\mathbf{p}_i	A feature vector of attributed sequence p_i .
μ_i	Attention weights.
α_i	Attention vector.

mation, and one Attention Block to learn the attention from both attributes and sequences.

6.2.1 Network Components.

6.2.1.1 AttNet

We build AttNet using fully connected neural network denoted as:

$$f(\mathbf{A}; \mathbf{W}_r, \mathbf{b}_r) = \tanh(\mathbf{W}_r \mathbf{A} + \mathbf{b}_r) \quad (6.2)$$

where \mathbf{W}_a and \mathbf{b}_a are two trainable parameters in the AttNet, denoting the weight matrix and bias vector, respectively. We use the activation

function `tanh` here in our `AttNet` based on our empirical studies. Other choices, such as `ReLU` or `sigmoid`, may work equally well in other real-world scenarios.

When given an attributed sequence $p_i = (\mathbf{v}_i, \mathbf{s}_i)$, `AttNet` takes the attributes as input and generates an attribute vector $\mathbf{r}_i = f(\mathbf{v}_i; \mathbf{W}_r, \mathbf{b}_r)$. Different from previous work in [10, 64] using stacked fully connected neural network as autoencoder, where the training goal is to minimize the reconstruction error, our goal of `AttNet` is to work together with other network components to maximize the possibility of predicting the correct labels.

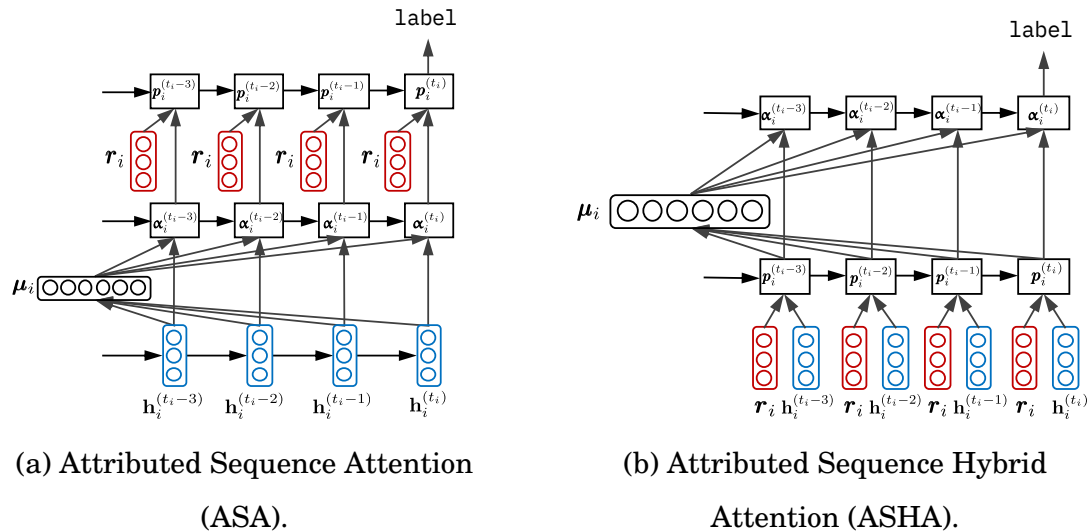


Figure 6.1: Two types of attention for attributed sequence classification in this task, where $\boldsymbol{\mu}_i = [\boldsymbol{\mu}_i^{(1)}, \dots, \boldsymbol{\mu}_i^{(t_i)}]$ is the attention weights.

6.2.1.2 SeqNet

Different from the attributes being unordered, items in our sequences have a temporal ordering. The information about the sequences is in both the item values and the ordering of items. The temporal orderings require a model that is capable of handling the dependencies between different items. There have been extensive studies on using recurrent neural networks (RNN) to handle temporal dependencies. However, RNN suffers from the problem of exploding and vanishing gradient during the training, where the gradient value becomes too large or too small and thus the network becomes untrainable. Long Short-Term Memory (LSTM) [25] is designed as one variation and expansion of the RNN to handle such issues. LSTM is capable of “remembering” values over long time intervals by introducing additional internal variables (*i.e.* various “gates” and “cell states”). We use an LSTM to handle the dependencies.

With a variable $\mathbf{X}^{(t)}$ at time t , the SeqNet can be expressed as:

$$\begin{aligned}
\mathbf{i}^{(t)} &= \sigma (\mathbf{W}_i \mathbf{X}^{(t)} + \mathbf{U}_i \mathbf{h}^{(t-1)} + \mathbf{b}_i) \\
\mathbf{f}^{(t)} &= \sigma (\mathbf{W}_f \mathbf{X}^{(t)} + \mathbf{U}_f \mathbf{h}^{(t-1)} + \mathbf{b}_f) \\
\mathbf{o}^{(t)} &= \sigma (\mathbf{W}_o \mathbf{X}^{(t)} + \mathbf{U}_o \mathbf{h}^{(t-1)} + \mathbf{b}_o) \\
\mathbf{g}^{(t)} &= \tanh (\mathbf{W}_c \mathbf{X}^{(t)} + \mathbf{U}_c \mathbf{h}^{(t-1)} + \mathbf{b}_c) \\
\mathbf{c}^{(t)} &= \mathbf{f}^{(t)} \odot \mathbf{c}^{(t-1)} + \mathbf{i}^{(t)} \odot \mathbf{g}^{(t)} \\
\mathbf{h}^{(t)} &= \mathbf{o}^{(t)} \odot \tanh (\mathbf{c}^{(t)})
\end{aligned} \tag{6.3}$$

where \odot denotes the bitwise multiplication, σ is a sigmoid activation function, $\mathbf{i}^{(t)}$, $\mathbf{f}^{(t)}$ and $\mathbf{o}^{(t)}$ are the internal gates of the LSTM, and $\mathbf{c}^{(t)}$ and $\mathbf{h}^{(t)}$ are the cell and hidden states of the LSTM, respectively. We denote the SeqNet as:

$$g(\mathbf{X}^{(t)}; \mathbf{W}_s, \mathbf{U}_s, \mathbf{b}_s) = \mathbf{h}^{(t)} \tag{6.4}$$

where $\mathbf{W}_s = [\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_c]$, $\mathbf{U}_s = [\mathbf{U}_i, \mathbf{U}_f, \mathbf{U}_o, \mathbf{U}_c]$ and $\mathbf{b}_s = [\mathbf{b}_i, \mathbf{b}_f, \mathbf{b}_o, \mathbf{b}_c]$.

With the sequence $\mathbf{s}_i = [\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(t_i)}]$ as part of an attributed sequence p_i , the hidden states for input $\mathbf{x}_i^{(t)}$ are $g(\mathbf{x}_i^{(t)}; \mathbf{W}_s, \mathbf{U}_s, \mathbf{b}_s) = \mathbf{h}_i^{(t)}$.

6.2.1.3 Attention Block

Recent work [45, 13] has identified that even LSTM-based solutions cannot fully handle the sequence learning on long sequences that the information over a long time may be lost. One popular solution to this problem is to incorporate the attention mechanism into the model. The attention mechanism effectively summarizes the data with the aim to leverage the importance of each item in the sequential input.

6.2.1.3.1 Attributed Sequence Attention (ASA). Different from the common sequence attention models, we now need to incorporate the attribute information into the learning process. Here, we design the Attention Block as follows: First, we need to compute the attention weight $\mu_i^{(t)}$ at t -th time as:

$$g(\mathbf{x}_i^{(t)}) = \mathbf{h}_i^{(t)}$$

$$\mu_i^{(t)} = \frac{\exp(g(\mathbf{x}_i^{(t)}))}{\sum_{j=1}^{t_i} g(\mathbf{x}_i^{(j)})}$$

Then, the attention weight is multiplied with the hidden state at each time step:

$$\alpha_i^{(t)} = \mu_i^{(t)} \odot g(\mathbf{x}_i^{(t)}), t = 1, 2, \dots, t_i \quad (6.5)$$

The attention weight $\mu_i^{(t)}$ at t time is randomly initialized and incrementally adjusted during the training process. The output at each time step is then augmented with the outputs from `AttNet` as:

$$\mathbf{p}_i^{(t)} = f(\mathbf{v}_i) \oplus \boldsymbol{\alpha}_i^{(t)}$$

At the last time step t_i , we denote $\mathbf{p}_i = \mathbf{p}_i^{(t_i)}$ to simplify the notation.

6.2.1.3.2 Attributed Sequence Hybrid Attention (ASHA). Different from the previous ASA approach, the outputs of `AttNet` and `SeqNet` are augmented with the `Attention Block`. The attention weight is written as:

$$d(\mathbf{v}_i, \mathbf{x}_i^{(t)}) = f(\mathbf{v}_i) \oplus g(\mathbf{x}_i^{(t)})$$

$$\boldsymbol{\mu}_i^{(t)} = \frac{\exp(d(\mathbf{v}_i, \mathbf{x}_i^{(t)}))}{\sum_{j=1}^{t_i} d(\mathbf{v}_i, \mathbf{x}_i^{(j)})}$$

Then, the vectors used for classification is:

$$\boldsymbol{\alpha}_i^{(t)} = \boldsymbol{\mu}_i^{(t)} \odot d(\mathbf{v}_i, \mathbf{x}_i^{(t)}), t = 1, 2, \dots, t_i \quad (6.6)$$

6.2.2 Attributed Sequence Classification

In the solution of attributed sequence classification *without* attention, the AttNet and the SeqNet are first concatenated as:

$$\mathbf{p}_i = d(\mathbf{v}_i, \mathbf{x}_i^{(t_i)}) = \mathbf{r}_i \oplus \mathbf{h}_i^{(t_i)} \quad (6.7)$$

Here, \oplus denotes the concatenation and t_i denotes the last item in s_i . Although all attributed sequences in the dataset are zero-padded to the maximum length t_{\max} , the padded zero values are masked and not used in the computation. We model the process of predicting the label for each attributed sequence as:

$$\Theta(p_i) = \begin{cases} \sigma(\mathbf{W}_p \mathbf{p}_i + \mathbf{b}_p), & \text{ASA or No Attention} \\ \sigma(\mathbf{W}_p \boldsymbol{\alpha}_i^{(t_i)} + \mathbf{b}_p), & \text{ASHA} \end{cases}$$

where σ is a sigmoid activation function and $\hat{l}_i = \Theta(p_i)$ is the predicted label. The \mathbf{W}_p and \mathbf{b}_p are both trainable in our model.

6.2.3 Training

6.2.3.1 Regularization

We adopt multiple strategies for different components in our AMAS network. We empirically select the following regularization strategies in our model based on: (1). For `SeqNet`, we apply ℓ_2 -regularization to the recurrent unit. (2). Dropout with a rate of 0.5 is used to regularize the fully connected layer in `AttNet`. (3). Lastly, we use Dropout with a rate of 0.2 in other fully connected layers in the model. Based on our observations, using regularization on `Attention Block` has no significant impact on the performance of AMAS.

6.2.3.2 Optimizer

We use an optimizer that computes the adaptive learning rates for every parameters, referred to as Adaptive Moment Estimation (`Adam`) [31]. The core idea is to keep (1). an exponentially decaying average of gradients in the past and (2). a squared past gradient. `Adam` counteracts the

biases as:

$$\widehat{\omega}^{(t)} = \frac{\beta_1 \omega^{(t-1)} + (1 - \beta_1) m^{(t)}}{1 - \beta_1^t}$$

$$\widehat{\nu}^{(t)} = \frac{\beta_2 \nu^{(t-1)} + (1 - \beta_2) (m^{(t)})^2}{1 - \beta_2^t}$$

where β_1 and β_2 are the decay rates, and $m^{(t)}$ is the gradient. We adopt $\beta_1 = 0.9$ and $\beta_2 = 0.999$ as in [31]. Finally, the Adam updates the parameters as:

$$\gamma^{(t+1)} = \gamma^{(t)} - \frac{\rho}{\sqrt{\widehat{\nu}^{(t)} + \epsilon}} \widehat{\omega}^{(t)}$$

where ρ is a static learning rate and ϵ is a constant with a small value to avoid division errors, such as division by zero. We empirically select $\rho = 0.01$.

6.3 Experiments

6.3.1 Datasets

Our solution has been motivated by use case scenarios observed at Amadeus corporation. For this reason, we work with the log files of an Amadeus [2] internal application. The log files contain user sessions in the form of at-

Table 6.2: Compared Methods

Name	Data Used	Attention	Note
BLA	Attributes	No	[24]
BLS	Sequences	No	[65]
BLAS	Attributes Sequences	No	[80]
SOA	Sequences	Yes	[76]
ASA	Attributes Sequences	Yes	This paper
ASHA	Attributes Sequences	Yes	This paper

tributed sequences. Also, we apply our methodology to real-world, publicly available Wikipedia data [68] and Reddit data [35]. For each type of data, we sample two subsets and conduct experiments independently. We summarize the data descriptions as follows:

- **Amadeus data (AMS-1, AMS-2)**¹. We sampled six datasets from the log files of an internal application at Amadeus IT Group. Each attributed sequence is composed of a user profile containing information (*e.g.* system configuration, office name) and a sequence of function names invoked by web click activities (*e.g.* login, search) ordered by time.
- **Wikipedia data (Wiki-1, Wiki-2)**. Wikipedia is an online game requiring participants to click through from a given start

¹Personal information is not collected.

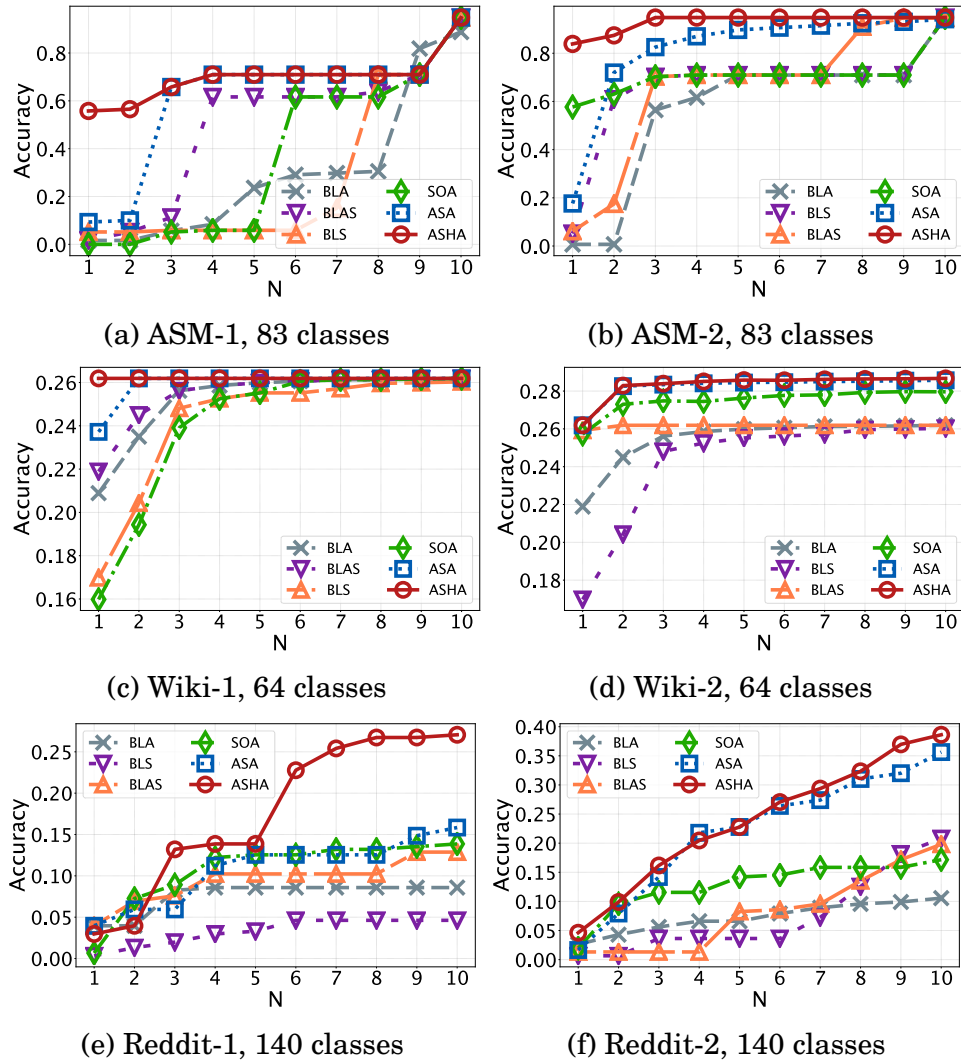


Figure 6.2: Performance comparison on all six datasets.

page to an end page using fewest clicks [68]. We select *finished* paths and extract several properties of each path (e.g., the category of the start path, time spent per click). We also sample six datasets

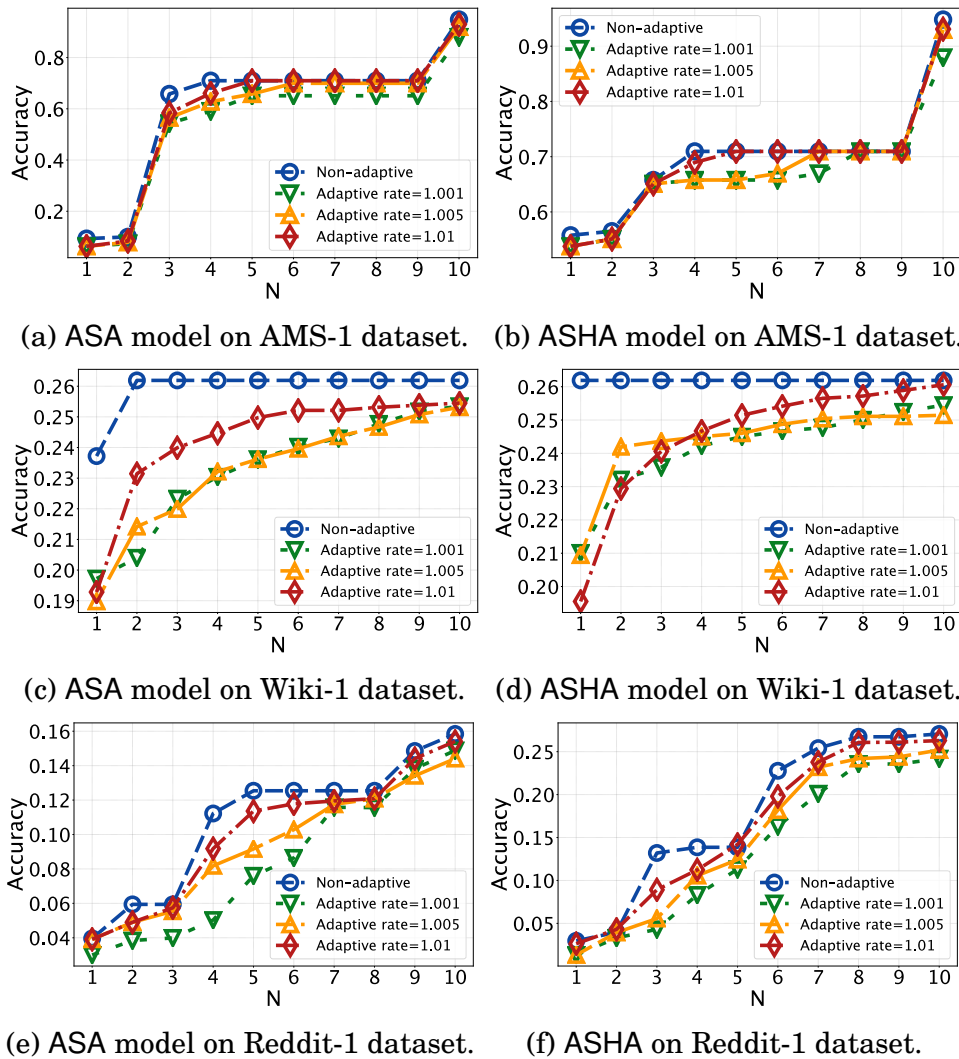


Figure 6.3: The performance comparison between non-adaptive and adaptive sampling.

from Wikipedia. The Wikipedia data is available through the Stanford Network Analysis Project² [37].

²<https://snap.stanford.edu/data/wikipedia.html>

- **Reddit data (Reddit-1, Reddit-2).** Reddit is an online forum. Two datasets that contain the content of reddit submissions are used. The Reddit data is available through the Stanford Network Analysis Project³.

We use 60% of the instances in each dataset for the training and the rest 40% for testing. In the training, we holdout 20% of the training instances for validation.

6.3.2 Compared Methods

We evaluate our two approaches, namely ASA and ASHA and compare them with the following baseline methods. We summarize all compared methods used in this research in Table [6.2](#).

- BLA is built using a fully connected neural network to reduce the dimensionality of the input data, and then classify each instance.
- BLS classifies sequences only data using an LSTM.
- BLAS utilizes the information from both attributes and sequences. The resulting embeddings generated by BLAS are then used for classification.

³<https://snap.stanford.edu/data/web-Reddit.html>

-
- SOA builds attention on the sequence data for classification, while the attribute data is not used.

6.3.3 Experimental Setting

Our paper focuses on the multi-class classification problem. We thus use accuracy as the metric to evaluate the performance. A higher accuracy score depicts more correct predictions of class labels. For each method, we holdout 20% as the validation dataset randomly selected from the training dataset. For each experimental setting, we report the top-1 ~ top-10 accuracy for each method.

We initialize our network using the following strategies: orthogonal matrices are used to initialize the recurrent weights, normalized random distribution [20] is used to initialize weight matrices in `AttNet`, and bias vectors are initialized as zero vector 0 .

6.3.4 Accuracy Results

In Figure [6.2], we compare the performance of our ASA and ASHA solutions with the other state-of-the-art methods in Table [6.2]. ASHA achieves the best performance of top-1 accuracy on most datasets. In most cases, ASHA outperforms other solutions significantly. We also observe a sig-

nificant performance improvement by ASA compared to other methods. That is, although the top-1 accuracy performance of ASA is beneath that of ASHA, it still outperforms SOA with sequence-only attention and all other methods without attention. The two closest competitors, the SOA utilizing the attention mechanism and classifying each instance based on only the sequential data, and BLAS using information from both attributes and sequences, but without the help from an attention mechanism, are outperformed by our proposed models.

6.3.5 Parameter Sensitivity Analysis

6.3.5.1 Adaptive Sampling Accuracy

As pointed out in recent work [10], adaptive sampling is capable of improving the efficiency of the optimization processes by adapting the training sample size in each iteration (*i.e.* epoch). In this set of experiments, we evaluate the two models with varying adaptive sampling rates. We use the adaptive sampling function as:

$$N_\tau = N_1 \lambda^{(\tau-1)}$$




Figure 6.4: Comparison of the history of training and validation losses. Here, τ denotes the epoch number, N_τ denotes the number of instances used in the τ -th epoch and λ the rate of adaptive sampling. We choose $\lambda = 1, 1.001, 1.005$, and 1.01 in our experiments, where $\lambda = 1$ means no adaptive sampling. The results presented in Figure 6.3 shows that the adaptive sampling with the above sampling rates can achieve similar performance as the non-adaptive approach yet now with much less training data.

adaptive training strategy. In Figure [6.4a](#), we observe the training with adaptive sampling is more aggressive compared to the non-adaptive approach. From Figure [6.4b](#) we conclude that with a higher adaptive rate, the model more easily becomes overfitted. Similar conclusion can also be made from Figure [6.4d](#). Selecting a higher adaptive sampling rate can shorten the training time but risking a higher chance of overfitting.

6.3.5.3 Case Studies

Figure [6.5](#) demonstrates the weights of each word often instances from the Reddit-2 dataset. Higher weights are represented with a darker color, while lower weights are represented with a lighter color. Comparing the three cases, we find that the SOA has the most polarized weights among the three cases. This may be caused by the fact that the attention produced by SOA is solely based on the sequences, while ASHA and ASA have been influenced by attribute data.

Chapter 7

Related Work

7.1 Deep Learning

Deep learning has received significant interests in various research areas in recent years. Deep learning models are capable of feature learning in varying granularities with hierarchical structures. Various deep learning models and optimization techniques have been proposed in a wide range of applications such as image recognition [28, 72] and sequence learning [12, 57, 73, 49]. Many of these applications involve the learning of single data type [12, 57, 73, 49], as other applications involve more than one data type [28, 72]. Several work [26, 46, 49] focuses on deep metric learning using deep learning techniques. The application of deep learning in sequence learning area has numerous work, one of the most popular work, sequence-to-sequence [57], uses long short-term memory model in machine translation. The hidden representations of sentences in the source language are transferred to a decoder to reconstruct in the target language. The idea is that the hidden representation can be used as a compact representation to transfer sequence similarities between two sequences. Multi-task learning [40] examines three multi-task learning settings for sequence-to-sequence models that aim at sharing either an encoder or decoder in an encoder-decoder model setting. Although the above work is capable of learning the dependen-

cies within a sequence, none of them focuses on learning the dependencies between attributes and sequences. Multimodal deep neural networks [28, 50, 72] is designed for information sharing across multiple neural networks, but none of these work focuses on our attributed sequence embedding problem.

7.2 One-shot Learning

One-shot learning has been known to be useful in various applications with very few training data [6, 59, 32]. The common problem setting of these tasks is to use one or few training examples to train a model that is capable of generalizing from the training examples and being used to predict the classes of previously unseen test data. The one-shot learning work dates back to work in [17] and [18], where the authors use a variational Bayesian framework for one-shot learning to categorize images by using few training examples. Recent work [6, 32] use discriminative network structure to train a learner for image classification tasks. Specifically, instead of using the siamese network for verification tasks, [32] uses the siamese network for classification tasks. Concerning the overfitting issue, [6] proposes an asymmetric variation of the siamese

network, named *learnnet*, to further address the issues of using a few training examples. Instead of using the siamese network, [59] focuses on using popular memory network and image attention to perform image classification tasks.

7.3 Attention Network

Attention network [45] has gained a lot research interest recently, the attention network has been applied in various tasks, including image captioning [72, 48], image generation [22], speech recognition [13] and document classification [76]. The goal of using attention network in these tasks is to make the neural network focus on the “interesting” parts of each input, such as, a small region of an image, or words that are helpful to classifying documents. There are different variations of attention network, including *hierarchical attention* [76] and *dual attentions* [48].

7.4 Sequence Mining

Recent work in sequence mining area aims at finding the most frequent subsequence pattern [44, 19]. Several recent work [4, 44] focus on finding the most frequent subsequence that meets certain constraints. That

is, find the set of sequential patterns satisfying various linguistic constraints (e.g., syntactic, symbolic). Many sequence mining work focuses on frequent sequence pattern mining. Recent work in [44] targets finding subsequences of possible non-consecutive actions constrained by a gap within sequences. [15] aims at solving pattern-based sequence classification problems using a parameter-free algorithm from the model space. It defines rule pattern models and a prior distribution on the model space. [19] builds a subsequence interleaving model for mining the most relevant sequential patterns. However, none of them learns sequence embeddings, nor do they support attribute data.

7.5 Clickstream Analysis

Recent studies [5, 36, 7] have studied clickstream analysis in various applications. Recent work [5] uses stream mining algorithms to identify frequent sequential patterns and then use the found patterns to build statistical models as Markov chains and transition matrices to capture frequent sequential patterns. In [7], the authors propose two frameworks using sequences to represent behaviors of students and then uses Markov chains to predict quiz performance. Targeting massive open on-

line courses, [36] and [7] focus on learning online student behaviors and performances based on clickstreams. [36] aims at finding the differences between the designed learning paths and the learning paths of students to improve the engagement and retention of students in online courses. [36] handles missing transitions between clicks and predict whether a sequence of clicks will result in an online purchasing. However, these work only focus on the clickstreams without the attribute information, nor do they generate embeddings.

7.6 Metric Learning

Distance metric learning, where the goal is to learn a distance metric from pairs of similar and dissimilar examples, has been studied in various work [70, 77, 14, 61, 43, 33, 26, 46, 49]. The common objective of these tasks is to learn a distance metric that the distance between similar pairs is reduced while the distance between dissimilar pairs is enlarged as much as possible. Distance metric learning has been used in various tasks to improve the performance of mining tasks, such as clustering [70, 77, 14]. Many applications in various domains also involve distance metric learning, including patient similarity in health in-

formatics [61], face verification in computer vision [43, 33, 26] and sentence semantic similarity analysis [46, 49]. However, these works only focus on the problem of metric learning on a single data type by using either attribute data or sequential data. Most of them focuses on linear transformation [61, 70, 77, 14].

Chapter 8

Conclusion and Future Work

8.1 Conclusion

The goal of this dissertation is to study the deep learning applications on attributed sequences. This dissertation studies four problems on the newly proposed attributed sequence data model, including attributed sequence embedding without labels, distance metric learning on attributed sequences, one-shot learning, and using attention model to filter useful information in the attributed sequence data. Here summarize the highlights of this dissertation:

First, we study the problem of *unsupervised attributed sequences embedding*. This work presents the design of the attributed sequence data

model and the design of the NAS framework. NAS is a novel deep learning-based framework that generates attributed sequence embeddings in an unsupervised setting. Different from conventional feature learning approaches, which work on either sequences or attributes without considering the *attribute-sequence dependencies*, we identify the three types of dependencies in attributed sequences. Our experiments on real-world tasks demonstrate that the proposed NAS effectively boosts the performance of outlier detection and clustering tasks compared to baseline methods.

Second, we target at incorporating the pairwise feedback into the embedding. In this work, we focus on the novel problem of distance metric learning on attributed sequences. We propose one MLAS with three solution variations to this problem using neural network models. The proposed MLAS network effectively learns the nonlinear distance metric from both attribute and sequence data, as well as the attribute-sequence dependencies. In our experiments on real-world datasets, we demonstrate the effectiveness of our MLAS network over other state-of-the-art methods in both performance evaluations and case studies.

Third, we study this new problem of one-shot learning for attributed sequences. We present the OLAS network design to tackle the chal-

lenges of utilizing this new data type in one-shot learning. OLAS incorporates two sub-networks, CoreNet , and PredictNet, that integrated into one structure together effectively learn the patterns hidden in this data type using only one example per class. OLAS uses this trained knowledge to generate labels for incoming unlabeled instances. Our experiments on real-world datasets demonstrate that OLAS on attributed sequences outperforms state-of-the-art one-shot learning methods.

Lastly, we propose a AMAS framework with two models for classifying attributed sequences. Our ASHA and ASA models progressively integrate the information from both attributes and sequences while weighing each item in the sequence to improve the classification accuracy. Experimental results demonstrate that our models significantly outperform state-of-the-art methods.

8.2 Future Work

The prevalence of attributed sequence data and the broad spectrum of real-world applications using attributed sequences motivate us to keep exploring this new direction of research. Based on the work in this dis-

sertation, there are several interesting research directions for future work:

First, in the attributed sequence embedding tasks, we currently only focus on using the information from attributes at the first time step. However, the information from attributes could be lost when encountered long sequences. One straightforward solution is to use attribute embedding to condition the sequence network at each step. Another interesting problem is to answer the question “*how will the embeddings be updated?*” That is, when the embedding model is updated, the embeddings before the update may be invalid. Other interesting directions include investigating how to building and use the index of a large number of attributed sequence embeddings.

Second, we now assume the feedback from domain experts is *triplets*, with two attributed sequences and one similarity label. However, feedback can be more complicated than pairwise attributed sequences. That is, domain experts may specify if a group of attributed sequences are similar or dissimilar. Another interesting direction would be using non-binary similarity label, where the similarity label now depicting the degrees of similarity of attributed sequences. There are also different ap-

proaches to build models for distance metric learning, such as matching network, worth investigating.

Third, there is a strong assumption in the one-shot learning task. That is, the unlabeled attributed sequence *must* belong to one of the known categories. However, this assumption may not always stand in real-world applications. For example, people will make fraud attempt that is novel to the model, in which case, the unlabeled attributed sequence may not have a match category. One-shot learning also assumes there is only one sample per category, and this problem can easily be extended to with *no* samples per category in real-world applications. For example, domain experts may know certain types of fraudulent transactions, but without usable cases. It would be interesting to investigate this direction as well.

Lastly, the goal of using the attention model in attributed sequence classification is to identify and give higher weights to the *relatively* useful items in the sequence. In addition to the two models proposed in this dissertation, one can augment the attribute network and sequence network with other layers to build attention model in various ways. It would be interesting to see the differences among different ways of building attention models.

Bibliography

- [1] Z. AKATA, F. PERRONNIN, Z. HARCHAOU, AND C. SCHMID, *Label-embedding for attribute-based classification*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013, pp. 819–826.
- [2] AMADEUS, *Amadeus IT Group*. <http://www.amadeus.com>. Accessed: 2017-09-23.
- [3] F. BASTIEN, P. LAMBLIN, R. PASCANU, J. BERGSTR, I. J. GOODFELLOW, A. BERGERON, N. BOUCHAR, AND Y. BENGIO, *Theano: new features and speed improvements*. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop, 2012.
- [4] N. BÉCHET, P. CELLIER, T. CHARNOIS, AND B. CRÉMILLEUX, *Sequence mining under multiple constraints*, in SIGSAC, 2015, pp. 908–914.
- [5] S. D. BERNHARD, C. K. LEUNG, V. J. REIMER, AND J. WESTLAKE, *Clickstream prediction using sequential stream mining techniques with markov chains*, in Proceedings of the 20th International Database Engineering & Applications Symposium, ACM, 2016, pp. 24–33.
- [6] L. BERTINETTO, J. F. HENRIQUES, J. VALMADRE, P. TORR, AND A. VEDALDI, *Learning feed-forward one-shot learners*, in NIPS, 2016, pp. 523–531.
- [7] C. G. BRINTON, S. BUCCAPATNAM, M. CHIANG, AND H. V. POOR, *Mining mooc clickstreams: Video-watching behavior vs. in-video quiz performance*, IEEE Transactions on Signal Processing, 64 (2016), pp. 3677–3692.

-
- [8] R. CAMPELLO, D. MOULAVI, A. ZIMEK, AND J. SANDER, *Hierarchical density estimates for data clustering, visualization, and outlier detection*, in Transactions on Knowledge Discovery from Data, ACM, 2015.
- [9] R. CHATPATANASIRI, T. KORSRILABUTR, P. TANGCHANACHAIANAN, AND B. KIJSIRIKUL, *A new kernelization framework for mahalanobis distance learning algorithms*, Neurocomputing, 73 (2010), pp. 1570–1579.
- [10] J. CHEN, S. SATHE, C. AGGARWAL, AND D. TURAGA, *Outlier detection with autoencoder ensembles*, in SDM, SIAM, 2017, pp. 90–98.
- [11] X.-W. CHEN AND X. LIN, *Big data deep learning: challenges and perspectives*, IEEE access, 2 (2014), pp. 514–525.
- [12] K. CHO, B. VAN MERRIËNBOER, C. GULCEHRE, D. BAHDANAU, F. BOUGARES, H. SCHWENK, AND Y. BENGIO, *Learning phrase representations using rnn encoder-decoder for statistical machine translation*, arXiv preprint arXiv:1406.1078, (2014).
- [13] J. K. CHOROWSKI, D. BAHDANAU, D. SERDYUK, K. CHO, AND Y. BENGIO, *Attention-based models for speech recognition*, in NIPS, 2015, pp. 577–585.
- [14] J. V. DAVIS, B. KULIS, P. JAIN, S. SRA, AND I. S. DHILLON, *Information-theoretic metric learning*, in International Conference on Machine Learning, 2007, pp. 209–216.
- [15] E. EGHO, D. GAY, M. BOULLÉ, N. VOISINE, AND F. CLÉROT, *A parameter-free approach for mining robust sequential classification rules*, in ICDM, 2015, pp. 745–750.
- [16] D. ERHAN, P.-A. MANZAGOL, Y. BENGIO, S. BENGIO, AND P. VINCENT, *The difficulty of training deep architectures and the effect of unsupervised pre-training*, in International Conference on Artificial Intelligence and Statistics, 2009, pp. 153–160.
- [17] L. FEI-FEI ET AL., *A bayesian approach to unsupervised one-shot learning of object categories*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2003, pp. 1134–1141.

-
- [18] L. FEI-FEI, R. FERGUS, AND P. PERONA, *One-shot learning of object categories*, IEEE transactions on pattern analysis and machine intelligence, 28 (2006), pp. 594–611.
- [19] J. FOWKES AND C. SUTTON, *A subsequence interleaving model for sequential pattern mining*, arXiv preprint arXiv:1602.05012, (2016).
- [20] X. GLOROT AND Y. BENGIO, *Understanding the difficulty of training deep feedforward neural networks*, in International Conference on Artificial Intelligence and Statistics, 2010, pp. 249–256.
- [21] A. GRAVES, *Generating sequences with recurrent neural networks*, arXiv preprint arXiv:1308.0850, (2013).
- [22] K. GREGOR, I. DANIHELKA, A. GRAVES, D. REZENDE, AND D. WIERSTRA, *Draw: A recurrent neural network for image generation*, in ICML, 2015, pp. 1462–1471.
- [23] R. HADSELL, S. CHOPRA, AND Y. LECUN, *Dimensionality reduction by learning an invariant mapping*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2006, pp. 1735–1742.
- [24] G. E. HINTON AND R. R. SALAKHUTDINOV, *Reducing the dimensionality of data with neural networks*, science, 313 (2006), pp. 504–507.
- [25] S. HOCHREITER AND J. SCHMIDHUBER, *Long short-term memory*, Neural computation, (1997).
- [26] J. HU, J. LU, AND Y.-P. TAN, *Discriminative deep metric learning for face verification in the wild*, in CVPR, 2014, pp. 1875–1882.
- [27] N. KALCHBRENNER AND P. BLUNSOM, *Recurrent continuous translation models.*, in Empirical Methods on Natural Language Processing (EMNLP), vol. 3, 2013, p. 413.
- [28] A. KARPATHY AND L. FEI-FEI, *Deep visual-semantic alignments for generating image descriptions*, in CVPR, 2015, pp. 3128–3137.
- [29] D. KEDEM, S. TYREE, F. SHA, G. R. LANCKRIET, AND K. Q. WEINBERGER, *Non-linear metric learning*, in NIPS, 2012, pp. 2573–2581.

-
- [30] A. KHALEGHI, D. RYABKO, J. MARI, AND P. PREUX, *Consistent algorithms for clustering time series*, Journal of Machine Learning Research, 17 (2016), pp. 1–32.
- [31] D. P. KINGMA AND J. L. BA, *Adam: A method for stochastic optimization*.
- [32] G. KOCH, R. ZEMEL, AND R. SALAKHUTDINOV, *Siamese neural networks for one-shot image recognition*, in ICML Deep Learning Workshop, 2015.
- [33] M. KOESTINGER, M. HIRZER, P. WOHLHART, P. M. ROTH, AND H. BISCHOF, *Large scale metric learning from equivalence constraints*, in CVPR, 2012, pp. 2288–2295.
- [34] B. M. LAKE, R. R. SALAKHUTDINOV, AND J. TENENBAUM, *One-shot learning by inverting a compositional causal process*, in NIPS, 2013, pp. 2526–2534.
- [35] H. LAKKARAJU, J. J. MCAULEY, AND J. LESKOVEC, *What’s in a name? understanding the interplay between titles, content, and communities in social media.*, ICWSM, 1 (2013), p. 3.
- [36] C. LAKSHMINARAYAN, R. KOSURU, AND M. HSU, *Modeling complex clickstream data by stochastic models: Theory and methods*, in WWW, 2016, pp. 879–884.
- [37] J. LESKOVEC, *Wikispeedia navigation paths*. <https://snap.stanford.edu/data/wikispeedia.html>. Accessed: 2018-04-09.
- [38] C.-Y. LIOU, W.-C. CHENG, J.-W. LIOU, AND D.-R. LIOU, *Autoencoder for words*, Neurocomputing, 139 (2014), pp. 84–96.
- [39] C.-Y. LIOU, J.-C. HUANG, AND W.-C. YANG, *Modeling word perception using the elman network*, Neurocomputing, 71 (2008), pp. 3150–3157.
- [40] M.-T. LUONG, Q. V. LE, I. SUTSKEVER, O. VINYALS, AND L. KAISER, *Multi-task sequence to sequence learning*, arXiv preprint arXiv:1511.06114, (2015).
- [41] L. V. D. MAATEN AND G. HINTON, *Visualizing data using t-sne*, Journal of Machine Learning Research, 9 (2008), pp. 2579–2605.

-
- [42] A. F. MCDAID, D. GREENE, AND N. HURLEY, *Normalized mutual information to evaluate overlapping community finding algorithms*, arXiv preprint arXiv:1110.2515, (2011).
- [43] A. MIGNON AND F. JURIE, *Pcca: A new approach for distance learning from sparse pairwise constraints*, in CVPR, 2012, pp. 2666–2672.
- [44] I. MILIARAKI, K. BERBERICH, R. GEMULLA, AND S. ZOUPANOS, *Mind the gap: Large-scale frequent sequence mining*, in SIGMOD, 2013, pp. 797–808.
- [45] V. MNIH, N. HEESS, A. GRAVES, ET AL., *Recurrent models of visual attention*, in NIPS, 2014.
- [46] J. MUELLER AND A. THYAGARAJAN, *Siamese recurrent architectures for learning sentence similarity.*, in AAAI, 2016, pp. 2786–2792.
- [47] V. NAIR AND G. E. HINTON, *Rectified linear units improve restricted boltzmann machines*, in International Conference on Machine Learning, 2010, pp. 807–814.
- [48] H. NAM, J.-W. HA, AND J. KIM, *Dual attention networks for multi-modal reasoning and matching*, (2017).
- [49] P. NECULOIU, M. VERSTEEGH, M. ROTARU, AND T. B. AMSTERDAM, *Learning text similarity with siamese recurrent networks*, Proceedings of the 1st Workshop on Representation Learning for NLP, (2016), p. 148.
- [50] J. NGIAM, A. KHOSLA, M. KIM, J. NAM, H. LEE, AND A. Y. NG, *Multimodal deep learning*, in ICML, 2011, pp. 689–696.
- [51] T. B. PEDERSEN AND C. S. JENSEN, *Multidimensional data modeling for complex data*, in Proceedings of the 15th IEEE International Conference on Data Engineering (ICDE), IEEE, 1999, pp. 336–345.
- [52] T. PHAM, T. TRAN, D. Q. PHUNG, AND S. VENKATESH, *Column networks for collective classification.*, in AAAI, 2017, pp. 2485–2491.
- [53] N. PHAN, Y. WANG, X. WU, AND D. DOU, *Differential privacy preservation for deep auto-encoders: an application of human behavior prediction.*, in AAAI, 2016, pp. 1309–1316.

-
- [54] A. SANTORO, S. BARTUNOV, M. BOTVINICK, D. WIERSTRA, AND T. LILICRAP, *Meta-learning with memory-augmented neural networks*, in International conference on machine learning, 2016, pp. 1842–1850.
- [55] A. M. SAXE, J. L. MCCLELLAND, AND S. GANGULI, *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*, arXiv preprint arXiv:1312.6120, (2013).
- [56] Y. SUN, Y. CHEN, X. WANG, AND X. TANG, *Deep learning face representation by joint identification-verification*, in NIPS, 2014, pp. 1988–1996.
- [57] I. SUTSKEVER, O. VINYALS, AND Q. V. LE, *Sequence to sequence learning with neural networks*, in NIPS, 2014, pp. 3104–3112.
- [58] A. TAJER, V. V. VEERAVALLI, AND H. V. POOR, *Outlying sequence detection in large data sets: A data-driven approach*, IEEE Signal Processing Magazine, 31 (2014), pp. 44–56.
- [59] O. VINYALS, C. BLUNDELL, T. LILICRAP, D. WIERSTRA, ET AL., *Matching networks for one shot learning*, in NIPS, 2016, pp. 3630–3638.
- [60] F. WANG AND J. SUN, *Survey on distance metric learning and dimensionality reduction in data mining*, Data Mining and Knowledge Discovery, 29 (2015), pp. 534–564.
- [61] F. WANG, J. SUN, AND S. EBADOLLAHI, *Integrating distance metrics learned from multiple experts and its application in patient similarity assessment*, in International Conference on Data Mining, SIAM, 2011, pp. 59–70.
- [62] G. WANG, X. ZHANG, S. TANG, H. ZHENG, AND B. Y. ZHAO, *Unsupervised clickstream clustering for user behavior analysis*, in Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, ACM, 2016, pp. 225–236.
- [63] J. WANG, A. KALOUSIS, AND A. WOZNICA, *Parametric local metric learning for nearest neighbor classification*, in NIPS, 2012, pp. 1601–1609.

-
- [64] W. WANG, Y. HUANG, Y. WANG, AND L. WANG, *Generalized autoencoder: A neural network framework for dimensionality reduction*, in CVPR, 2014, pp. 490–497.
- [65] Y. WANG AND F. TIAN, *Recurrent residual learning for sequence classification*, in EMNLP, 2016, pp. 938–943.
- [66] W. WEI, J. LI, L. CAO, Y. OU, AND J. CHEN, *Effective detection of sophisticated online banking fraud on extremely imbalanced data*, World Wide Web, 16 (2013), pp. 449–475.
- [67] R. WEST AND J. LESKOVEC, *Human wayfinding in information networks*, in Proceedings of the 21st International Conference on World Wide Web, ACM, 2012, pp. 619–628.
- [68] R. WEST, J. PINEAU, AND D. PRECUP, *Wikispeedia: An online game for inferring semantic distances between concepts*, in IJCAI, 2009.
- [69] WIKIPEDIA, *Internet bot — wikipedia, the free encyclopedia*. https://en.wikipedia.org/wiki/Internet_bot, 2017. Accessed: 2017-10-25.
- [70] E. P. XING, M. I. JORDAN, S. J. RUSSELL, AND A. Y. NG, *Distance metric learning with application to clustering with side-information*, in NIPS, 2003, pp. 521–528.
- [71] Z. XING, J. PEI, AND E. KEOGH, *A brief survey on sequence classification*, SIGKDD, 12 (2010), pp. 40–48.
- [72] K. XU, J. BA, R. KIROS, K. CHO, A. COURVILLE, R. SALAKHUDINOV, R. ZEMEL, AND Y. BENGIO, *Show, attend and tell: Neural image caption generation with visual attention*, in ICML, 2015, pp. 2048–2057.
- [73] Y. XU, J. H. LAU, T. BALDWIN, AND T. COHN, *Decoupling encoder and decoder networks for abstractive document summarization*, MultiLing 2017, (2017), p. 7.
- [74] R. R. YAGER AND N. ALAJLAN, *Probabilistically weighted owa aggregation*, IEEE Transactions on Fuzzy Systems, 22 (2014), pp. 46–56.

-
- [75] J. YANG AND W. WANG, *Cluseq: Efficient and effective sequence clustering*, in Proceedings of the 19th IEEE International Conference on Data Engineering (ICDE), IEEE, 2003, pp. 101–112.
- [76] Z. YANG, D. YANG, C. DYER, X. HE, A. J. SMOLA, AND E. H. HOVY, *Hierarchical attention networks for document classification*, in HLT-NAACL, 2016, pp. 1480–1489.
- [77] D.-Y. YEUNG AND H. CHANG, *A kernel approach for semisupervised metric learning*, IEEE Transactions on Neural Networks, 18 (2007), pp. 141–149.
- [78] P. ZAVLARIS, *If you build it, they will come.* <https://resources.distilnetworks.com/all-blog-posts/bad-bot-report-2017-intro>. Accessed: 2017-10-25.
- [79] Z. ZHUANG, X. KONG, E. RUNDENSTEINER, J. ZOUAOUI, AND A. ARORA, *Unsupervised attributed sequence embedding: a deep learning approach*, In submission. <http://goo.gl/hfnSQB>.
- [80] Z. ZHUANG, X. KONG, E. RUNDENSTEINER, J. ZOUAOUI, AND A. ARORA, *Unsupervised embedding of attributed sequences: A multimodal learning approach*, in submission, 2018.