

Goal-Driven Context-Aware Service Recommendation for Mashup Development

Xihao Xie, Jia Zhang
Department of Computer Science
Southern Methodist University
 Dallas, USA
 {xihaox; jiazhang}@smu.edu

Rahul Ramachandran
Marshall Space Flight Center
 NASA
 Huntsville, USA
 rahul.ramachandran@nasa.gov

Tsengdar J. Lee
Science Mission Directorate
NASA Headquarters
 Washington, USA
 tsengdar.j.lee@nasa.gov

Seungwon Lee
Jet Propulsion Laboratory
 NASA
 Pasadena, USA
 seungwon.lee@jpl.nasa.gov

Abstract—As service-oriented architecture becoming one prevalent technique to rapidly compose functionalities to customers, increasingly more reusable software components have been published online in the form of web services. To create a mashup, however, it gets not only time-consuming but also error-prone for developers to find suitable services components from such a sea of services. Service discovery and recommendation has thus attracted significant momentum in both academia and industry. This paper proposes a novel incremental *recommend-as-you-go* approach to recommending next potential service based on the context of a mashup under construction, considering services that have been selected up to the current step as well as the mashup goal. The core technique is an algorithm of learning the embedding of services, which learns their past goal-driven context-aware decision making behaviors in addition to their semantic descriptions and co-occurrence history. A goal exclusionary negative sampling mechanism tailored for mashup development is also developed to improve training performance. Extensive experiments on a real-world dataset demonstrate the effectiveness of this approach.

Index Terms—Service recommendation, Mashup creation, Service embedding

I. INTRODUCTION

In the last two decades, a huge number of software components have been published onto the Internet in the form of software services (or so-called APIs or services in short). Up to June 15th, 2022, ProgrammableWeb¹, the largest online repository of APIs, has accumulated 24,471 services. Such remotely accessible services enable software developers to compose existing services into mashups, easier and faster than before. However, such a sea of services makes it not only time-consuming but also error-prone for mashup developers to pick up suitable service candidates. Thus, service recommendation-powered mashup development has attracted significant momentum in recent years.

Fig. 1 illustrates a possible development process for a mashup that we randomly selected from ProgrammableWeb named *Shared Count*². As shown in Fig. 1, a mashup creation is an incremental process with multiple rounds of service selection, e.g., *Twitter* is selected first followed by *Facebook*. In each round, an instant recommender suggests the “next” suitable services.

¹<https://www.programmableweb.com/>

²<https://www.programmableweb.com/mashup/shared-count>

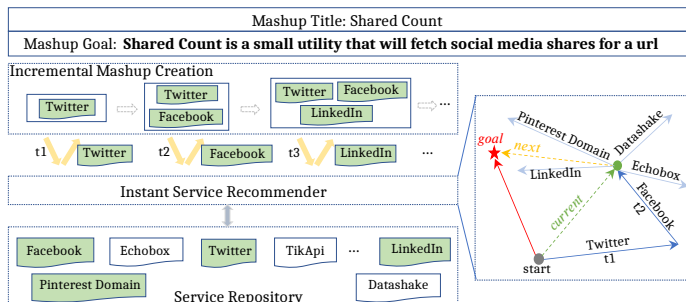


Fig. 1. Incremental process of mashup creation.

In our position, an effective recommender shall guide a mashup development process as a journey of exploration in a forest (i.e., an embedding space), led by a light beam (i.e., the mashup goal), as illustrated in the right-hand side dotted box in Fig. 1. At a point of the decision making process, e.g., the green dot, the services that have already been selected (i.e., *Twitter* at step t_1 and *Facebook* at step t_2) represent the position of the “traveler” at the time. The problem thus becomes how to make the next step (i.e., select the next service) toward the ultimate goal (i.e., the red star representing the embedding of the textual mashup goal “shared count is a small utility that will fetch social media shares for a url”). We coin a term *context* to represent a decision making point during a mashup development process, which comprises selected services and the mashup goal each being denoted as a *contextual item* throughout this paper.

To tackle this decision making problem, this paper proposes a goal-driven, context-aware machine learning method capable of recommending “next” suitable services in each step of the incremental process of mashup creation. To the best of our knowledge, we make the first effort to learning the embedding of goal-driven mashup context. For each service, in addition to learning its semantic descriptions and service co-occurrence, we learn its decision making behaviors from past mashup development provenance. Furthermore, our proposed learning framework is extensible for considering other dimensions of data such as user profile. The main contributions of our work are three-fold. First, we propose a novel machine learning algorithm that is capable of learning goal-driven context

embedding during mashup development. Second, we propose a goal-exclusionary negative sampling strategy tailored for a rapid training process for mashup development. Third, we demonstrate the effectiveness of our approach based on extensive experiments on a real-world dataset.

The remainder of this paper is organized as follows. Section II summarizes the related work. Section III defines the problem and presents our approach in detail. Section IV explains how to train and learn relevant parameters. Section V presents and analyzes experimental results. Finally, Section VI draws conclusions.

II. RELATED WORK

In this section, we will discuss closely related work from two aspects: service representation learning and next service recommendation.

A. Service Representation Learning

In recent years, researchers have explored many ways to represent services in a vector space to support downstream service recommendation applications. Here are some examples. Gu et al. [2] propose an approach to recommending a bundle of services by leveraging latent Dirichlet allocation (LDA) to embed services and mashups in a topic space. Wang et al. [3] propose Service2vec to learn service representations based on a constructed service network. Zhang et al. [4] incorporate service users' perceptions into service profiles to form more comprehensive service representations. Menzi et al. [5] embed services into a low-dimensional vector space based on a constructed context-aware service knowledge graph to support service recommendation.

In contrast to their work, our approach learns service embeddings from past service selection decision making behaviors, in addition to service and mashup descriptions and correlations.

B. Next Service Recommendation

Next service recommendation aims to recommend candidate services for developers in each step of a multi-round process of mashup creation. Cao et al. [1] design a collaborative-filtering (CF)-based algorithm based on a two-level topic model to recommend services. Zhang et al. [6] propose an approach to extracting relationships of people, services and workflows from historical usage data into a social network to proactively recommend services in a workflow composition process. Liu et al. [7] develop a generalized sequential pattern algorithm to mine frequent composition patterns of mashups, and design an interactive recommendation algorithm to assist mashup creation. By modeling the relations between services and service-based systems (SBSs) into a heterogeneous information network, Xie et al. [8] measure semantic similarities between SBSs and use content filtering technique to recommend next service. Kirubananthan et al. [14] propose a method to support long-term service composition recommendation according to user ratings.

Our work differs from the aforementioned studies in two aspects. First, we introduce a model to embed the context

including not only existing services but also mashup goals. In this way, our model can learn goal-driven next service prediction behaviors from past experiences. Second, we model different contributions of different contextual items with an attention mechanism.

III. CONTEXT-AWARE NEXT SERVICE RECOMMENDATION

Fig. 2 depicts an overview of our goal-driven context-aware service recommendation framework. Given a partial mashup under construction, all services already selected and the mashup goal form the context of the recommendation problem.

As shown in Fig. 2, for each contextual service item, its semantic embedding and auxiliary embedding (to be learned) are fused into an intermediate vector. Based on the services selected, an attention function (to be learned) will generate attention coefficients for them. All weighted service embeddings will be fused and, then fused with the intermediate vector from the semantic embedding and auxiliary embedding (to be learned) of the mashup goal. A fully connected layer with a weight matrix (to be learned) will compute the possibility of each service becoming the next service, and make the prediction via a softmax function.

Take the motivating mashup in Fig. 1 again as an example, with m being its textually described goal. At its second creation step, two services have been selected as contextual services, s_i being *Twitter* and s_k being *Facebook*. After training, the framework in Fig. 2 will recommend a candidate service, s_n being *LinkedIn*, for the third step.

In this section, we first define the research problem and introduce some preliminary, and then discuss in detail our approach.

A. Research Problem Statement and Preliminary

Our study is based on a repository comprising all historical mashups and their component services. Let $\mathcal{M} = \{m_1, m_2, \dots, m_{|\mathcal{M}|}\}$ be a set of mashups, each of which is textually described in natural language as its goal, and $\mathcal{S} = \{s_1, s_2, \dots, s_{|\mathcal{S}|}\}$ be a set of services. Each mashup m_i contains a subset of services $S^{m_i} = \{s_1^i, s_2^i, \dots, s_{|S^{m_i}|}^i\} \subseteq \mathcal{S}$. In an incremental process of creating mashup $m' \notin \mathcal{M}$, before it is completed, given multiple services that have been selected as contextual services $S_t^{m'} = \{s_1^{m'}, s_2^{m'}, \dots, s_t^{m'}\} \subseteq \mathcal{S}$, where $s_j^{m'} \in \mathcal{S}, j \in [1, t]$ is the service selected at the j^{th} step, our goal is to calculate the probability for each service $s \in \mathcal{S} \setminus S_t^{m'}$ and recommend top K candidate services at the next step $t + 1$. Further, let $C_t^{m'} = S_t^{m'} \cup \{m'\}$ denote the whole context at the t^{th} step of creating mashup m' . In this way, the recommendation problem can be formulated to a problem of probabilistic classification by learning to predict a conditional probability:

$$p(s_{t+1}^{m'} = s | C_t^{m'}) = p(s | C_t^{m'}) \quad (1)$$

To solve the research problem, we first introduce some preliminary.

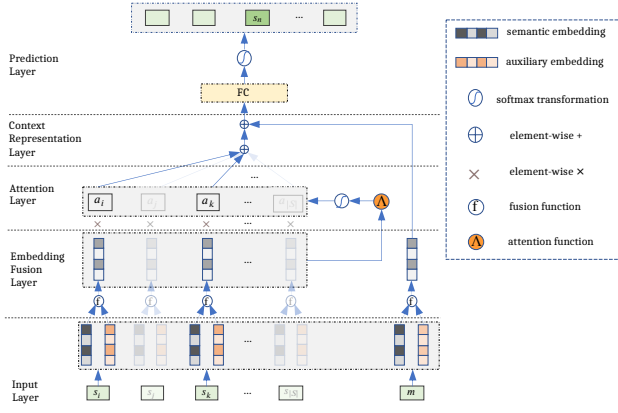


Fig. 2. Blueprint of proposed approach.

a) *Semantic Embedding Matrix*: Let matrix $\mathbf{E}^S \in \mathbb{R}^{(|S|+1) \times d}$ be a semantic vector space for services and mashup goal. The j^{th} row of the matrix \mathbf{E}_j^S , $j \in [1, |S|]$, is a d -dimension vector semantically representing the j^{th} service $s_j \in \mathcal{S}$. The last row of \mathbf{E}^S , i.e., $\mathbf{E}_{|S|+1}^S$, is the semantic embedding for mashup goal. This matrix aims to learn from textual descriptions of all services and mashups from the repository.

b) *Auxiliary Embedding Matrix*: Let matrix $\mathbf{E}^X \in \mathbb{R}^{(|S|+1) \times d}$ denote an auxiliary vector space for services and mashup goal. The j^{th} row of the matrix \mathbf{E}_j^X , $j \in [1, |S|]$, is a d -dimension vector representation for the j^{th} service $s_j \in \mathcal{S}$. The last row of \mathbf{E}^X , i.e., $\mathbf{E}_{|S|+1}^X$ is the auxiliary embedding for mashup goal. This matrix aims to learn from past service occurrence and decision making behaviors.

c) *Attention Vector*: Let $\mathbf{A} \in \mathbb{R}^d$ be a global d -dimension vector that is shared by all services. It is used to model the contribution scales of contextual services based on an attention mechanism.

d) *Weight Matrix*: Let $\mathbf{E}^W \in \mathbb{R}^{d \times |S|}$ be a weight matrix of the fully connected layer, as shown in Fig. 2. It is used to help predict the conditional probability defined in Eq. (1).

In our framework, $\Theta = \{\mathbf{E}^X, \mathbf{A}, \mathbf{E}^W\}$ will be learned through back-propagation, which will be discussed in detail in the following sections.

B. Context Embedding

The semantic embedding matrix \mathbf{E}^S is learned through NLP machine learning methods such as the Doc2Vec [13] model. Each textual description of mashup goal or service is regarded as a document. The corpus of all documents is fed into the Doc2Vec model to receive a vectorized representation of each mashup goal or service. Specially, we adopt the distributed memory model of paragraph vectors (PV-DM) as it performs better than the distributed bag of words of paragraph vector (PV-DBOW) [13]. The trained \mathbf{E}^S will also be used for goal exclusionary negative sampling, which will be discussed later in Section 4. The auxiliary embedding matrix \mathbf{E}^X carries non-semantic information, e.g., co-occurrence information and decision-making strategies, of services and mashup. For any

two services items for example, the closer they are in \mathbf{E}^X , the more likely they are to co-occur in the same context. Note that $\mathbf{E}_{|S|+1}^X$, the last row of the matrix, is an in-situ vector that will be shared by all mashups. The vector will be trained as an offset vector to adjust the semantic representation of a mashup goal in a decision making process.

Given a partial mashup $C_t^{m'}$, the input units in the input layer of Fig. 2 constitute a one-hot encoding vector where only the unit at position i ($s_i \in C_t^{m'}$) is set to 1. Thus, for each contextual service $s_i \in C_t^{m'}$, the corresponding i^{th} rows of the trained \mathbf{E}^S and \mathbf{E}^X , $i \in [1, |S|]$, are fed into the fusion layer. For the mashup goal, its semantic embedding $\mathbf{E}_{|S|+1}^S$ will be obtained in an *ad hoc* manner from the trained Doc2Vec model, and its auxiliary embedding will be obtained from the trained $\mathbf{E}_{|S|+1}^X$.

As shown in Fig. 2, for each i^{th} contextual item, its semantic embedding and auxiliary embedding are fused into a fused vector in the fusion layer: $\mathbf{E}_i^F = f(\mathbf{E}_i^S, \mathbf{E}_i^X)$, where $i \in [1, |S| + 1]$, and $f: \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}^d$ is a fusion function.

C. Attention-based Context Embedding

Considering that different contextual services may weight differently in representing the whole context, we model the contribution scale of every contextual service with an attention mechanism. Specifically, for each contextual service $s_{t,k}^{m'} \in C_t^{m'}$, its contribution scale is denoted as the weight of $w_{t,k}^{m'}$, where $k \in [1, |S|]$. It can be calculated with a softmax transformation which is widely used in neural networks:

$$w_{t,k}^{m'} = \frac{\exp(\Lambda(k))}{\sum_{s_{t,h}^{m'} \in C_t^{m'}} \exp(\Lambda(h))} \quad (2)$$

$$\Lambda(k) = \mathbf{E}_k^F \cdot \mathbf{A} = f(\mathbf{E}_k^S, \mathbf{E}_k^X) \cdot \mathbf{A} \quad (3)$$

where $\Lambda(k)$ is an attention function shown in Fig. 2. It is a dot product of the fused vector \mathbf{E}_k^F with an attention vector \mathbf{A} to be learned, which is used to model the contribution scales of contextual services.

The resulted integration weight $w_{t,k}^{m'}$ acts as the attention scale of contextual service $s_{t,k}^{m'}$. The higher $w_{t,k}^{m'}$ is, the more we should pay attention to the contextual service $s_{t,k}^{m'}$ than other services $s_{t,l}^{m'} \in C_t^{m'} \setminus \{s_{t,k}^{m'}\}$ at the t^{th} step of creating mashup m' .

As shown in Fig. 2, we integrate the representations of all contextual services and that of mashup goal to represent the whole context $C_t^{m'}$ as follows:

$$\mathbf{E}^{C_t^{m'}} = \mathbf{E}_{|S|+1}^F + \sum_{s_{t,k}^{m'} \in C_t^{m'}} w_{t,k}^{m'} \times \mathbf{E}_k^F, \quad s.t. \quad \sum_{s_{t,k}^{m'} \in C_t^{m'}} w_{t,k}^{m'} = 1 \quad (4)$$

D. Prediction & Recommendation

Since the context $C_t^{m'}$ has been represented as Eq. 4, the conditional probability of Eq. 1 becomes:

$$p(s_{t+1}^{m'} = s | C_t^{m'}) = p(s | \mathbf{E}^{C_t^{m'}}) \quad (5)$$

As shown in Fig. 2, in the fully connected layer, the j^{th} column of the weight matrix (to be learned) \mathbf{E}_j^W is used to calculate a scalar value to help obtain the conditional probability. Specifically, for every service $s_n \in \mathcal{S}$, the scalar value can be calculated as a dot production of the context embedding $\mathbf{E}_{t'}^{C_t^{m'}}$ and \mathbf{E}_n^W :

$$v(n) = \mathbf{E}_{t'}^{C_t^{m'}} \cdot \mathbf{E}_n^W \quad (6)$$

Here $v(n)$ can be seen as a scoring function qualifying the relevance of a service s_n with respect to the given context $C_t^{m'}$. For a service s_n , the higher the score is, the more relevant it is to the context, and the more likely it should be recommended. Thereafter, for any service $s_n \in \mathcal{S}$, the likelihood of it to be selected at the next step defined in Eq. 5 can be further normalized as a result of the softmax transformation:

$$p_{\Theta}(s_{t+1}^{m'} = s_n | \mathbf{E}_{t'}^{C_t^{m'}}) = \frac{\exp(v(n))}{\sum_{s_{t,h}^{m'} \in \mathcal{S}} \exp(v(h))} \quad (7)$$

Now that the probability of a service to be selected at the next step has been calculated, for all of the services in $\mathcal{S} \setminus S_t^{m'}$ that have not been selected, we rank them according to their probabilities in descending order, and recommend the top K services.

In practice, at the beginning of creating a mashup, only the mashup goal will be used as the input of our model. That is, according to Eq. (4), the embedding of context becomes: $\mathbf{E}_{0'}^{C_0^{m'}} = \mathbf{E}_{|\mathcal{S}|+1}^F = f(\mathbf{E}_{|\mathcal{S}|+1}^S, \mathbf{E}_{|\mathcal{S}|+1}^X)$.

IV. MODEL LEARNING

In this section, we explain how to train our model, i.e., $\Theta = \{\mathbf{E}^X, \mathbf{A}, \mathbf{E}^W\}$. The matrices are learned throughout creation steps of mashups from historical data offline.

A. Objective Function

Let $\mathcal{D} = \{D^i\}$ denote a training data set, where $D^i = \bigcup_{0 \leq t < |S^{m_i}|} \{D_t^i\}$; $D_t^i = \langle C_t^i, G_t^i \rangle$ is a training instance from mashup m_i at the t^{th} step of creating it; $C_t^i = \{m_i\} \cup S_t^{m_i}$ is the given context, and $G_t^i \subseteq \mathcal{S} \setminus S_t^{m_i}$ is the set of the observed ground truth services that co-occur with $S_t^{m_i}$ in the mashup at step t . Given a training set \mathcal{D} , the joint probability distribution can be obtained as:

$$p_{\Theta}(\mathcal{D}) \propto \prod_{D^i \in \mathcal{D}} \prod_{D_t^i \in D^i} p_{\Theta}(G_t^i | \mathbf{E}_{t'}^{C_t^i}) \quad (8)$$

Thereafter, parameters Θ can be learned by maximizing the following objective function over all training instances:

$$\begin{aligned} L_{\Theta} &= \log \prod_{D^i \in \mathcal{D}} \prod_{D_t^i \in D^i} \sum_{s_{t,j}^i \in G_t^i} p_{\Theta}(s_{t,j}^i | \mathbf{E}_{t'}^{C_t^i}) \\ &= \log \prod_{D^i \in \mathcal{D}} \prod_{D_t^i \in D^i} \frac{\sum_{s_{t,j}^i \in G_t^i} \exp(v(j))}{\sum_{s_h \in \mathcal{S}} \exp(v(h))} \end{aligned} \quad (9)$$

where $p_{\Theta}(s_{t,j}^i | \mathbf{E}_{t'}^{C_t^i})$ is the conditional probability according to Eq. 7.

B. Sampling Strategy

Optimizing the objective function in Eq. 9, however, is non-trivial since each evaluation of the softmax function has to traverse all services, which is extremely time-consuming. To learn the aforementioned parameters efficiently, we employ the idea of negative sampling [9] to approximate the conditional probability. Considering that a service being negative to a mashup m_i tends to be less compatible to m_i , we introduce a goal exclusionary negative sampling strategy to obtain a set of negative services. Specifically, based on the trained Doc2Vec model, for all services in $\mathcal{S} \setminus S^{m_i}$, we rank them according to their cosine-based semantic similarities with the mashup goal in ascending order and choose the top r percent as negative samples. In this way, let $N^i \subseteq \mathcal{S} \setminus S^{m_i}$ denote the set of negatively sampled services, each of which is not contained in the mashup m_i . Then, we redefine the conditional probability as:

$$p_{\Theta}(G_t^i, N^i | \mathbf{E}_{t'}^{C_t^i}) = \prod_{s_h \in G_t^i \cup N^i} p_{\Theta}(s_h | \mathbf{E}_{t'}^{C_t^i}) \quad (10)$$

$$p_{\Theta}(s_h | \mathbf{E}_{t'}^{C_t^i}) = \begin{cases} \sigma(v_{\Theta}(h)) & s_h \in G_t^i \\ 1 - \sigma(v_{\Theta}(h)) & s_h \in N^i \end{cases} \quad (11)$$

where σ is the sigmoid function $\sigma(x) = 1/(1 + \exp(-x))$ and $v_{\Theta}(h)$ is the scoring function following the definition of Eq. 6. In Eq. 11, $\sigma(v_{\Theta}(h))$ can be seen as the probability of service s_h being labeled as a ground truth service with a score of $v_{\Theta}(h)$ given context of C_t^i . Thus, the probability of it being labeled as a negative service is $1 - \sigma(v_{\Theta}(h))$. In this way, maximizing the conditional probability in Eq. 10 means maximizing the probability of a service being a positive sample and in the meantime minimizing the probability of it being a negative sample. After that, the objective function becomes:

$$\begin{aligned} L_{\Theta} &= \log \prod_{D^i \in \mathcal{D}} \prod_{D_t^i \in D^i} \prod_{s_h \in G_t^i \cup N^i} p_{\Theta}(s_h | \mathbf{E}_{t'}^{C_t^i}) \\ &= \sum_{D^i \in \mathcal{D}} \sum_{D_t^i \in D^i} \left\{ \sum_{s_h \in G_t^i} \log \sigma(v_{\Theta}(h)) + \sum_{s_h \in N^i} \log[\sigma(-v_{\Theta}(h))] \right\} \end{aligned} \quad (12)$$

Let $L_{\Theta}^{i,t}$ denote the objective function for a training instance D_t^i as:

$$L_{\Theta}^{i,t} = \sum_{s_h \in G_t^i} \log \sigma(v_{\Theta}(h)) + \sum_{s_h \in N^i} \log[\sigma(-v_{\Theta}(h))] \quad (13)$$

At last, according to Eq. 12, we utilize the stochastic gradient ascent to learn the aforementioned parameters through back-propagation.

The gradient and update rule of \mathbf{E}_h^W then can be represented as:

$$\frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{E}_h^W} = [l(h) - \sigma(v_{\Theta}(h))] \cdot \mathbf{E}_{t'}^{C_t^i}, \mathbf{E}_h^W \leftarrow \mathbf{E}_h^W + \eta \cdot \frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{E}_h^W} \quad (14)$$

where $l(h)$ is a label function whose value is 1 when $s_h \in G_t^i$ or 0 otherwise. The gradient and update rule of \mathbf{A} for D_t^i can be represented as:

$$\frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{A}} = [l(h) - \sigma(v_{\Theta}(h))] \cdot \mathbf{E}_h^W \cdot \frac{\partial \mathbf{E}^{C_t^i}}{\partial \mathbf{A}}, \mathbf{A} \leftarrow \mathbf{A} + \eta \cdot \frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{A}} \quad (15)$$

Similarly, the gradient and update rule of \mathbf{E}_k^X for $s_k \in S_t^{m_i}$ and $k = |S| + 1$ becomes:

$$\frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{E}_k^X} = [l(h) - \sigma(v_{\Theta}(h))] \cdot \mathbf{E}_h^W \cdot \frac{\partial \mathbf{E}^{C_t^i}}{\partial \mathbf{E}_k^X}, \mathbf{E}_k^X \leftarrow \mathbf{E}_k^X + \eta \cdot \frac{\partial L_{\Theta}^{i,t}}{\partial \mathbf{E}_k^X} \quad (16)$$

For the auxiliary embedding of mashup goal, $\frac{\partial \mathbf{E}^{C_t^i}}{\partial \mathbf{E}_{|S|+1}^X} = 1 \cdot \frac{\partial \mathbf{E}_{|S|+1}^F}{\partial \mathbf{E}_{|S|+1}^X}$.

C. Model Training

Algorithm 1 Learning parameters \mathbf{E}^X , \mathbf{A} and \mathbf{E}^W offline

Input: mashups \mathcal{M} , services \mathcal{S} in training set, dimension size d , learning rate η and negative sampling ratio r

Output: parameters \mathbf{E}^X , \mathbf{A} and \mathbf{E}^W

```

1: Initialize  $\mathbf{E}^X \in \mathbb{R}^{(|S|+1) \times d}$ ,  $\mathbf{A} \in \mathbb{R}^d$  and  $\mathbf{E}^W \in \mathbb{R}^{d \times |S|}$ 
2:  $\mathbf{E}_{:|S|}^S \leftarrow PV - DM(\mathcal{M}, \mathcal{S}, d)$  // obtain semantic embedding
   of services from PV-DM
3:  $L \leftarrow 0$ 
4: while  $L$  not converged do
5:   for each mashup  $m_i \in \mathcal{M}$  do
6:      $\mathbf{E}_{|S|+1}^S \leftarrow PV - DM(m_i)$ 
       // obtain semantic embedding of mashup goal from PV-
       DM
7:      $N^i \leftarrow NegativeSampling(\mathbf{E}^S, r)$  // Sample top  $r$ 
       percent negative services
8:     for each creation step  $t$  of  $m_i$  do
9:        $D_t^i \leftarrow GenerateInstances(m_i, S)$  // generate  $D_t^i$  as
        $\langle C_t^i, G_t^i \rangle$ 
10:       $\mathbf{E}^{C_t^i} \leftarrow EmbedContext(\mathbf{E}^S, \mathbf{E}^X, \mathbf{A}, C_t^i)$  // context
       embedding using Eq. (4)
11:       $L^{i,t} \leftarrow ObjectiveFunctionValue(\mathbf{E}^W, \mathbf{E}^{C_t^i}, G_t^i, N^i)$ 
       // Compute value of objective function using Eq. (13)
12:       $UpdateParameters(G_t^i, N^i, \mathbf{E}^{C_t^i}, \eta)$ 
       // Update parameters according to the rules in Eq. (14),
       (15) and (16)
13:       $L \leftarrow L + L^{i,t}$ 
14:     end for
15:   end for
16: end while
17: return  $\mathbf{E}^X, \mathbf{A}, \mathbf{E}^W$ 

```

Algorithm 1 lists the pseudo code of model training. The first two steps are preprocessing: we first initialize parameters (\mathbf{E}^X , \mathbf{A} and \mathbf{E}^W) randomly; and then feed the corpus of textual descriptions of mashup goals and services into the PV-DM model [13] to obtain semantic embedding of services, i.e., $\mathbf{E}_{:|S|}^S$. Afterwards, for each mashup m_i in the training set (line 5), we obtain its semantic embedding (line 6), i.e., $\mathbf{E}_{|S|+1}^S$, from the PV-DM model, and sample the top r percent goal exclusionary negative services (line 7). After that, incrementally, at every creation step t (line 8), we generate a training instance D_t^i (line 9), containing contextual services $S_t^{m_i}$, mashup goal, and ground truth services G_t^i . Specifically,

$S_t^{m_i}$ is augmented by randomly choosing one service from G_{t-1}^i . Subsequently, G_t^i becomes $S^{m_i} \setminus S_t^{m_i}$. Afterwards, we embed the context as a d -dimension vector according to Eq. (4) (line 10). We then compute the objective function value according to Eq. (13) (line 11), and update \mathbf{E}^X , \mathbf{A} and \mathbf{E}^W (line 12) according to Eqs. (14), (15) and (16) respectively. Finally, we stop learning if the accumulated objective function values converge.

Note that our offline training stage is goal driven in two aspects. First, the auxiliary embedding of mashup goal and contextual services, i.e., $\mathbf{E}_{|S|+1}^X$ and $\mathbf{E}_{:|S|}^X$, are jointly and reciprocally learned based on the semantic embedding of mashup goal, i.e., $\mathbf{E}_{|S|+1}^S$. Second, in the negative sampling phase, mashup goal is leveraged to sample negative services that are semantically exclusionary to it.

V. EXPERIMENTS

In this section, we explain experiments conducted and analyze the results.

A. Experimental Setup

We designed two research questions to evaluate our approach from two angles:

- **RQ1:** How does our approach perform in terms of recommendation metrics?
- **RQ2:** How does our approach perform throughout the incremental process?
- **RQ3:** How does the attention mechanism affect the performance?

We evaluated our approach on a real-world dataset that was crawled from ProgrammableWeb up to December 2020. The mashups containing only one service were removed. Consequently, we obtained 1,553 mashups and 663 services to form our testbed. Table I summarizes the overall statistics of our dataset.

TABLE I
STATISTIC INFORMATION OF THE DATASET

Statistics	Values
# of Mashups	1,553
# of Services	663
Avg. # of services per mashup	2.70
# of Training instances	3,310
# of Testing instances	883

During the training stage, we randomly chose 80% of mashups as the training set and the rest 20% ones as the test set. Without losing generality, we built training instances at the last step of mashup construction. For each mashup m_i in the training set, we constructed $|S^{m_i}|$ training instances. In detail, for mashup m_i , each service $s_k^i \in S^{m_i}$ was used as ground truth to be recommended, and the rest services of S^{m_i} were treated as the corresponding contextual services. That is to say, D^i was expanded as $D^i = \cup_{s_k^i \in m_i} \{D_k^i\}$, where $D_k^i = \langle C^i(k), s_k^i, N^i \rangle$ and $C^i(k) = \{m_i\} \cup S^{m_i} \setminus \{s_k^i\}$ is the corresponding context of s_k^i .

For textual descriptions of services and mashup goals, we performed a series of pre-processing operations including tokenization, stop-words removal and lemmatization with the NLTK³ library. After that, we utilized the PV-DM model from the doc2vec library of Gensim⁴ to obtain semantic vectorized embeddings (i.e., \mathbf{E}^S). As for N^i , we set the negative sampling ratio (i.e., r) to be 20%.

We adopted the *leave-last-out* evaluation scheme in the experiments, meaning that we evaluated recommendation performance at the last step of constructing the testing mashups. Similar to the expansion operation adopted for the training set, for each mashup m_j in the testing set, we expanded it to $|S^{m_j}|$ testing instances. For each $s_k^j \in S^{m_j}$, it was used as the ground truth service and the remainder of S^{m_j} , $C^j(k) = \{m_j\} \cup S^{m_j} \setminus \{s_k^j\}$, was used as corresponding context.

Four competitive methods were used as the baselines:

- **MatchUp** [10]: A service recommender system which is based on correlations of used services as the user context.
- **SNRec** [11]: An approach to recommending next service according service similarities based on co-existence.
- **WVSM** [12]: A content-based method to recommend services using textual content similarity.
- **ISRec** [8]: A clustering and classification-based service recommendation approach with word embeddings.

Note that, in our test bed, services that are published after the mashup were not available to be recommended. Thus, for all methods, the services whose submission time is earlier than that of the mashup were held out for recommendation. In addition, for simplicity, we used element-wise plus as the fusion function: $\mathbf{E}^F = f(\mathbf{E}^S, \mathbf{E}^X) = \mathbf{E}^S \oplus \mathbf{E}^X$. Thus, in Eq. (16), $\frac{\partial \mathbf{E}^{C_i}}{\partial \mathbf{E}_X^{S_{i+1}}} = 1$.

B. Evaluation Metrics

We employed the following two commonly used accuracy metrics for evaluation. For both metrics, the higher the value is, the better the performance is.

1. **REC@K**: It measures the recall of the top K recommended services over all testing set. We reported the results with $K \in \{3, 5, 10, 20\}$.

$$REC@K = \frac{|R^K \cap G|}{|G|} \quad (17)$$

where R^K and G are the recommended list of services and the ground truth services, respectively.

2. **MRR**: It measures the reciprocal rank of the position of the ground truth over all testing set.

$$MRR = \frac{1}{rank(R^K, G)} \quad (18)$$

where $rank(R^K, G)$ is the hit position of the ground truth in the recommended result.

³<http://www.nltk.org/>

⁴<https://radimrehurek.com/gensim/models/doc2vec.html>

C. Performance Comparison

For MatchUp, we defined the importance of a service as the number of mashups in which it is invoked [10]. The importance of a mashup was defined as the mean importance of its component services. For WVSM, being the same with [12], we set the smooth efficient λ to 0.1. For ISRec, except for the number of topics T , which was set to be 28, we set other parameters the same with [8]. For our approach, the dimension size d and learning rate η were set to be 100 and 0.001, respectively. Table II shows the results of mean metric values of different methods over all testing instances. The best is in bold and the second best is underlined. In terms of all accuracy metrics, the experiment demonstrates that our approach outperforms others, which answers the first research question **RQ1**.

TABLE II
PERFORMANCE COMPARISON OF DIFFERENT APPROACHES

Methods	REC@3	REC@5	REC@10	REC@20	MRR
MatchUp	0.2514	0.2899	0.3386	0.3613	0.2099
SNRec	0.2661	0.3205	0.3692	0.4145	0.1738
WVSM	0.2276	0.2820	0.3533	0.4723	0.1790
ISRec	<u>0.3228</u>	<u>0.4009</u>	<u>0.4881</u>	<u>0.5946</u>	<u>0.2325</u>
Ours	0.3579	0.4602	0.5277	0.6116	0.2872

Both MatchUp and SNRec take into account invocations between mashups and services. However, they do not consider mashup goal and the contribution scales of contextual services. The reasons that our approach performs better than WVSM are two-fold. First, our approach takes into account not only textual semantic information, but also decision-making information. Second, we take contributions of different contextual services into consideration. As for ISRec, our approach performs better because it does not make use of the different contribution scales of selected services.

Especially, in terms of MRR, the results of other approaches are lower than 0.25. However, our approach reaches at 0.2872, meaning that generally the hit service will be ranked at the third or forth position of the recommended result, which is quite useful to support mashup development.

D. Recommendation Performance in Mashup Creation Process

Ideally, the recommendation of our approach shall help mashup development incrementally approaching mashup goal. In other words, as the creation step increases, the recommendation results of our approach should become increasingly better. In our test bed, we noticed that the number of creation steps (i.e., the number of composing services) differs in different mashups. Thus, to make it easier to answer **RQ2**, we divided the mashups into groups, each containing mashups with the same number of creation steps. For example, g_3 groups the mashups comprising three services. For a mashup m_j in group g_l , we measured the recommendation performance over the number of creation step t ranging from 0 to $l - 1$. At creation

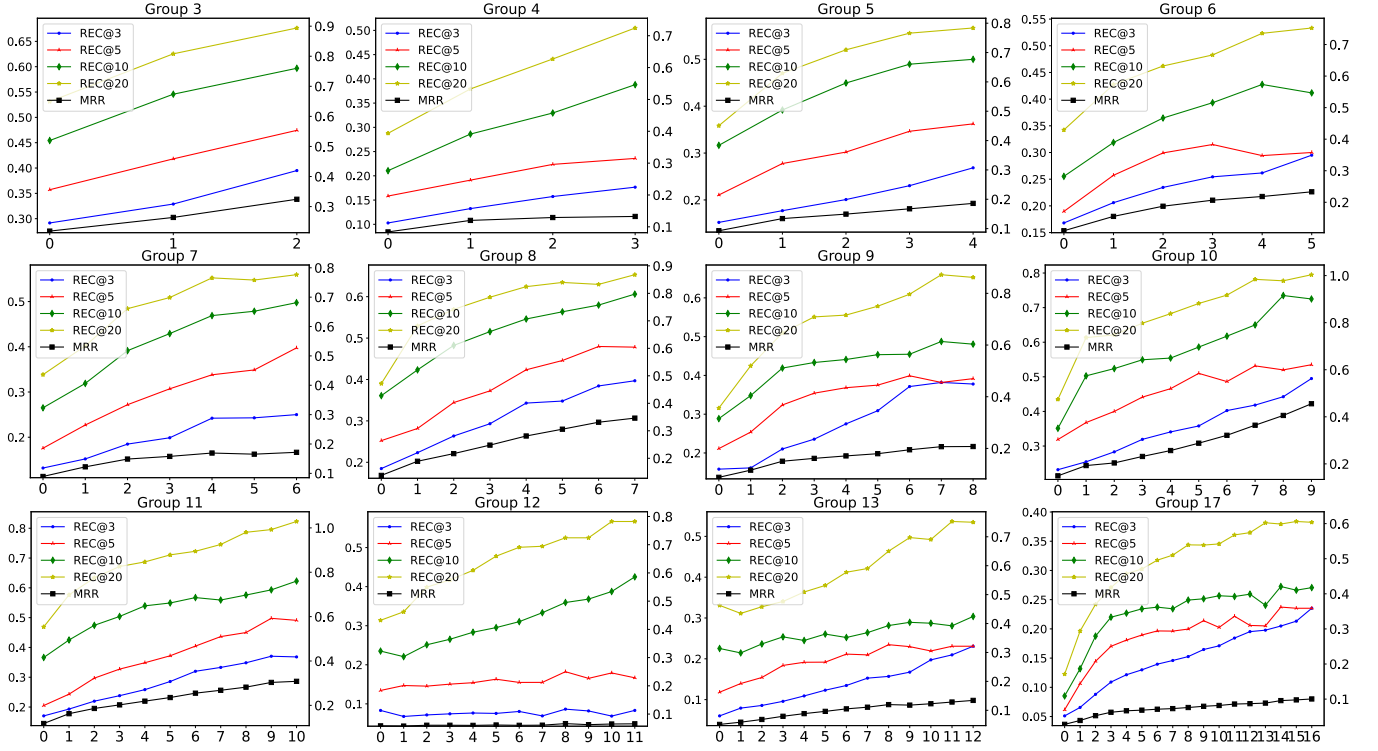


Fig. 3. Recommendation result in terms of accuracy metrics in mashup creation process

step t , the contextual services S_t^j and ground truth services G_t^j are generated as we described previously in Section IV-C.

Fig. 3 presents the average values of accuracy metrics over each number of creation step t in different mashup groups $g_3 \sim g_{17}$. For each sub-figure in Fig. 3, x-axis represents the index of creation step (t). The left y-axis and right y-axis are mean values of evaluation metrics of REC@K and MRR defined in Eqs. (17) and (18), respectively. The black line depicts the average values of MRR over all test cases (i.e., use the right y-axis). The other colored lines show the test results of REC@K when K=3,5,10,20, respectively.

As shown in Fig. 3, for almost all testing groups, the recommendation performance goes up consistently as t increases. In some groups, though, the accuracy metrics sometimes go down. The reason we found is that some services are only included in the mashups in the testing set. Thus, the embeddings of those services are never learned in the training phase. Examples are two out of eight services, 66,366⁵ and 148,157⁶, in the mashup 190,280⁷ in group 8.

Overall, considering the average ratio of such “long-tail” services over all mashups is lower than 8.23%, the experimental results align with our intuition: with the number of steps t getting higher, the more information we could extract from the context, the better the recommendation performance could be, and the closer to mashup goal the recommended result is.

E. Effect of Attention Mechanism

In this subsection, we will answer the third research question **RQ3** by investigating the effect of the attention mechanism in our approach. Fig. 4 shows the effect of the attention mechanism in terms of the accuracy metrics of REC@K and MRR. In terms of REC@3, REC@5, REC@10 and REC@20, the recommendation performance increases by 6.91%, 6.69%, 6.28% and 5.71%, respectively. In terms of MRR, we observe a similar result, as the performance increases by 3.79%. The results are in line with the empirical fact that adopting the attention mechanism could enhance the expression ability of the embeddings of services and finally improve the overall performance.

Fig. 5 illustrates the visualization of attention weights on six example mashups and the corresponding 20 services. The horizontal and vertical coordinates are IDs of mashups and services in ProgrammableWeb, respectively. In mashup 87,810⁸, two services contribute the most: a chrome extension pulling social network sharing counts into the chrome browser, and a service 68,272⁹ for accessing detailed user data on Pinterest. This experiment shows that, given a set of selected services in the incremental mashup creation process, the attention mechanism is *adaptive*, *expressive* and *context-aware* to represent the latent requirement of the unfinished mashup.

⁵<https://www.programmableweb.com/api/skimlinks-link-monetization>

⁶<https://www.programmableweb.com/api/parsebot>

⁷<https://www.programmableweb.com/mashup/christmas-list-app>

⁸<https://www.programmableweb.com/mashup/sharemetric-chrome-extension>

⁹<https://www.programmableweb.com/api/pinterest-domain>

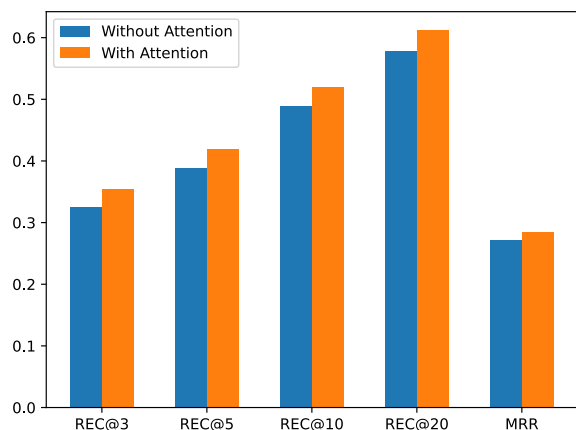


Fig. 4. The effect of attention mechanism

VI. CONCLUSIONS

We regard mashup creation as an incremental process which at each step, we can help developers speed up the mashup development by recommending the next service. In this paper, we have presented a novel machine learning method of goal-driven context embedding, which learns decision making strategies for service composition toward mashup goals in addition to service co-occurrence and semantic descriptions. To efficiently train parameters, we have introduced a goal exclusionary negative sampling strategy. Note that our approach is general and can be expanded to broader next candidate recommendation scenarios that are goal oriented, such as next product recommendation in e-shopping systems.

In the future, we plan to extend our research in three directions. First, we argue that a developer's preference plays an important role in her mashup creation process, thus we plan to take personalized information into account to further enhance our approach. Second, our current work applies a naïve function, i.e., element-wise plus, for embedding fusion. We plan to explore other fusion functions to make our model more expressive. Third, we observed in the incremental process of mashup creation, the order of contextual services is an influencing factor for next service selection. We will take the sequential creation behavior into consideration to further consolidate our approach.

REFERENCES

- [1] Cao, B., Liu, X. F., Rahman, M. M., Li, B., Liu, J., Tang, M.: Integrated Content and Network-Based Service Clustering and Web APIs Recommendation for Mashup Development. *IEEE Transactions on Services Computing* **13**(1), 99–113 (2017)
- [2] Gu, Q., Cao, J., Liu, Y.: CSBR: A Compositional Semantics-based Service Bundle Recommendation Approach for Mashup Development. *IEEE Transactions on Services Computing*, doi: 10.1109/TSC.2021.3085491 (2021)
- [3] Wang, S., Wang, Z., Xu, X.: Mining Bilateral Patterns as Priori Knowledge for Efficient Service Composition. In *Proceedings of IEEE International Conference on Web Services*, 65–72. San Francisco, CA, USA (2016)

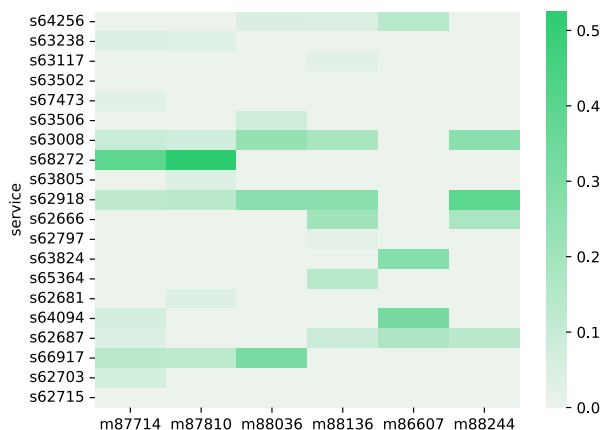


Fig. 5. Visualization of attention weights

- [4] Zhang, J., Fan, Y., Zhang, J., Bai, B.: Learning to Build Accurate Service Representations and Visualization. *IEEE Transactions on Services Computing*, **15**(3), 2022, 1551–1563.
- [5] Mezni, H., Benslimane, D., Belletreche, L.: Context-aware Service Recommendation based on Knowledge Graph Embedding. *IEEE Transactions on Knowledge and Data Engineering*, doi: 10.1109/TKDE.2021.3059506 (2021)
- [6] Zhang, J., Tan, W., Alexander, J., Foster, I., Madduri, R.: Recommend-As-You-Go: A Novel Approach Supporting Services-oriented Scientific Workflow Reuse. In *Proceedings of IEEE International Conference on Services Computing*, 48–55. Washington DC, USA (2011)
- [7] Liu, X., Sun, H., Wu, H., Zhang, R., Liu, X.: Using Sequential Pattern Mining and Interactive Recommendation to Assist Pipe-like Mashup Development. In *Proceedings of 8th IEEE International Symposium on Service Oriented System Engineering*, 173–180. Oxford, UK (2014)
- [8] Xie, F., Wang, J., Xiong, R., Zhang, N., Ma, Y., He, K.: An Integrated Service Recommendation Approach for Service-based System Development. *Expert Systems with Applications* **123**, 178–194 (2019)
- [9] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., Dean, J.: Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of The 26th International Conference on Neural Information Processing Systems* **2**, 3111–3119 (2013)
- [10] Abiteboul, S., Greenshpan, O., Milo, T., Polyzotis, N.: MatchUp: Autocompletion for Mashups. In *Proceedings of IEEE International Conference on Data Engineering*, 1479–1482. Shanghai, China (2009)
- [11] Pan, W., Li, B., Shao, B., He, P.: Service Classification and Recommendation Based on Software Networks. *Chinese Journal of Computers* **34**(12), 2355–2369 (2011)
- [12] Li, C., Zhang, R., Huai, J., Sun, H.: A Novel Approach for API Recommendation in Mashup Development. In *Proceedings of IEEE International Conference on Web Services*, 289–296. Anchorage, AK, USA (2014)
- [13] Le, Q., Mikolov, T.: Distributed Representations of Sentences and Documents. In *Proceedings of IEEE International Conference on Machine Learning*, 1188–1196. Anchorage, AK, USA (2014)
- [14] Kirubanantham, P., Sankar, S.M.U., Amuthadevi, C., Baskar, M., Raja, M.S., Karthik, P.C.: An Intelligent Web Service Group-based Recommendation System for Long-term Composition. *The Journal of Supercomputing*, **78**, 1944–1960 (2022)