

A Machine Learning Approach for Automated Filling of Categorical Fields in Data Entry Forms

HICHEM BELGACEM, University of Luxembourg, Luxembourg

XIAOCHEN LI*, Dalian University of Technology, China

DOMENICO BIANCULLI, University of Luxembourg, Luxembourg

LIONEL BRIAND, University of Luxembourg, Luxembourg and University of Ottawa, Canada

Users frequently interact with software systems through data entry forms. However, form filling is time-consuming and error-prone. Although several techniques have been proposed to auto-complete or pre-fill fields in the forms, they provide limited support to help users fill categorical fields, i.e., fields that require users to choose the right value among a large set of options.

In this paper, we propose LAFF, a learning-based automated approach for filling categorical fields in data entry forms. LAFF first builds Bayesian Network models by learning field dependencies from a set of historical input instances, representing the values of the fields that have been filled in the past. To improve its learning ability, LAFF uses local modeling to effectively mine the local dependencies of fields in a cluster of input instances. During the form filling phase, LAFF uses such models to predict possible values of a target field, based on the values in the already-filled fields of the form and their dependencies; the predicted values (endorsed based on field dependencies and prediction confidence) are then provided to the end-user as a list of suggestions.

We evaluated LAFF by assessing its effectiveness and efficiency in form filling on two datasets, one of them proprietary from the banking domain. Experimental results show that LAFF is able to provide accurate suggestions with a Mean Reciprocal Rank value above 0.73. Furthermore, LAFF is efficient, requiring at most 317 ms per suggestion.

CCS Concepts: • **Computing methodologies** → *Bayesian network models*; • **Information systems** → **Recommender systems**; • **Software and its engineering** → **Software usability**.

Additional Key Words and Phrases: Form filling, Data entry forms, Machine Learning, Software data quality, User interfaces

ACM Reference Format:

Hichem Belgacem, Xiaochen Li, Domenico Bianculli, and Lionel Briand. 2022. A Machine Learning Approach for Automated Filling of Categorical Fields in Data Entry Forms. *ACM Trans. Softw. Eng. Methodol.* 1, 1, Article 1 (January 2022), 40 pages. <https://doi.org/10.1145/3533021>

1 INTRODUCTION

Data entry forms are essential software user interface (UI) elements [42, 68] to collect users' inputs. According to statistics, approximately 70 million professionals or 59% of all professionals in the United States need to complete on-line forms for their daily jobs [77]. However, form filling is time-consuming and error-prone [64]. Some domain-specific software, such as enterprise resource

*This work was done while the author was affiliated with the University of Luxembourg, Luxembourg.

Authors' addresses: Hichem Belgacem, University of Luxembourg, Luxembourg, hichem.belgacem@uni.lu; Xiaochen Li, Dalian University of Technology, China, xiaochen.li@dlut.edu.cn; Domenico Bianculli, University of Luxembourg, Luxembourg, domenico.bianculli@uni.lu; Lionel Briand, University of Luxembourg, Luxembourg, University of Ottawa, Canada, lionel.briand@uni.lu.

© 2022 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Software Engineering and Methodology*, <https://doi.org/10.1145/3533021>.

planning (ERP) systems [2], may include data entry forms with hundreds of fields for users to fill in during the execution of specific business processes.

The form filling process lays an extra burden on users, who have to spend tremendous energy on inputting the right value into each field, resulting in a task that is both slow and frustrating [33]. This situation inevitably leads to data quality issues (e.g., when the wrong value is provided as input for a field), which may seriously affect data-reliant software systems and undermine business opportunities and even cause loss of human life [5]. Existing statistics reveal that data entry typically has an error rate of over 1% [24]. Such erroneous data is then transferred into the software system that uses it and affects all the other connected information. For example, erroneous input data in retailing systems alone leads to a waste of \$2.5 billion each year for consumers [24]. Further, data entry errors in spreadsheets mislead business decisions, causing additional costs in correcting the errors [65]. Even worse, data entry has been a top cause of medication errors [5], which resulted in at least 24 deaths in the US in 2003 [57].

From a software engineering point of view, it is important to develop approaches that reduce data entry errors. Since data entry forms may contain various types of fields (e.g., textual, numerical, and categorical fields), various studies have proposed different strategies to improve data quality. One way is to detect data quality issues by running test queries (which check the semantic validity of the data) on the application database right after data updates; this is the approach proposed in the context of continuous data testing [57], which focuses on identifying numerical errors. Another possible solution for data error prevention is to relieve users from the burden of form filling, i.e., automating the form filling process, with mechanisms that pre-fill or auto-complete data entry forms during data entry sessions [66, 73]. Existing work in this area mainly focuses on textual fields; they build language models (e.g., sequence-to-sequence models) to learn the character or word relation from historical textual inputs, and then provide word auto-completion based on the letters typed in a field [82].

However, the aforementioned techniques cannot be applied to *categorical* fields, which are fields that provide a list of candidate values (also called “options”) from which the user has to choose (e.g., country). Such fields are likely to generate data quality issues. For instance, empirical studies have shown that more than half (54.5%) of the data errors in a medical record system were caused by the candidate value selection error [61]. Users could wrongly select the job role, the modality of care, or the drugs unintentionally [49, 61], causing potential medical negligence. As another example, in financial systems, the correctness of selected values is required by regulations (e.g., for the “purpose of business” or “source of wealth” fields, when opening a bank account), to prevent money laundering [11].

The data entry errors in categorical fields are mainly caused by two reasons. First, categorical fields require focused attention for users to choose among a large set of options in a limited time. For example, a real-world computer-assisted clinical ordering system contains a lengthy alphabetized list of 33 items, 13 of which are category headers with subsumed (concealed) levels [40]. Also, a species management system may require users to select species from dozens of genera, with nearly a hundred valid species in a single genus [28]. It is well known that users tend to make more mistakes as more options are presented to select from [17, 49]. Moreover, juxtaposition selection errors can happen when adjacent values are textually similar (e.g., in an alphabetically-sorted list) [70]. Although some systems enable searchable list boxes to help users accelerate this selection [43], filling categorical fields is still difficult, as users require significant cognitive effort to match each option with the actual value they intend to fill in. Many users lack a detailed conceptual model of the software system [40] defined by requirements analysts and domain experts. For example, junior healthcare providers may not remember and understand all the options predefined in the field of “modality of care” in an electronic medical record system [61], which necessitates a potential

lengthy search process or may lead to the selection of an inappropriate value [40]. The cognitive load is even higher when there are more options for users to compare with [17]. In such cases, the time to select a value is linear with the list length [20].

Some form filling tools [32, 77] support filling categorical fields, by suggesting frequently selected values in a field or values selected by a user in some similar fields from 3rd-party software systems. Nevertheless, they provide limited support due to the low accuracy of their suggestions; moreover, their usage may violate enterprise security policies, since they rely on information from 3rd-party software systems [79].

Furthermore, existing automated form filling approaches exhibit some limitations when dealing with (1) forms filled following an arbitrary order, and (2) partially filled forms. The first case occurs because, while filling in a data entry form, it is very frequent to have little or no restriction on the order of users' inputs: a user may select any field as the next target or even go back and modify already filled fields. Some automated form filling approaches [38] require a fixed form filling order before building a recommendation model; however, this assumption is unrealistic from a practical standpoint. As for the second case, at a certain time during the data entry session, a form is usually partially filled, meaning that an approach for automated form filling can only use information in currently filled fields. However, when this is not sufficient to predict the target, existing approaches [38, 53] yield inaccurate suggestions.

To provide effective form filling suggestions, in this paper, we propose *LAFF*, a Learning-based Automated Form Filling approach, for filling categorical fields. The basic idea of *LAFF* is to build machine learning models based on input instances (i.e., fields and the corresponding values provided in input) obtained from data entry forms that have been filled in the past (hereafter called *historical input instances*); such models represent dependencies among fields in historical input instances. Using these models, the already-filled fields in a data entry form can then be used as features to predict the possible values of a given target field. *LAFF* aims to be used by developers, who can integrate it into their data entry form implementations.

To deal with forms filled in an arbitrary order (which would result in a huge number of combinations of filled fields (features) and target fields to handle), *LAFF* utilizes Bayesian Networks (BN) to mine the dependencies between any field combinations, without assuming, a priori, an order for form filling. Moreover, to improve its learning ability, *LAFF* uses a local modeling strategy to cluster historical input instances; further, it builds additional local BNs, which learn fine-grained field dependencies from the clusters of historical input instances. These local models capture additional dependencies that might not have been captured by the model trained on the entire historical dataset. Once the trained models are available, *LAFF* uses them in the form filling suggestion phase, which occurs during the data entry session: given a target field, *LAFF* selects one of the available BNs and predicts the possible values of the field based on the values in the already-filled fields. To deal with partially filled forms (which might lead to inaccurate suggestions), *LAFF* includes a heuristic-based endorser, which determines whether the values predicted in the previous step are accurate enough to be returned to the user, based on the analysis of the field dependencies and of the predicted probability distribution of the values for the target field.

We evaluated *LAFF* using form filling records from both a public dataset and a proprietary dataset extracted from a production-grade enterprise information system in the banking domain. The experimental results show *LAFF* can yield a large number of accurate suggestions with a *Mean Reciprocal Rank (MRR)* value above 0.73 and a prediction coverage rate ranging from 0.70 to 0.87, significantly outperforming a state-of-the-art approach based on association rule mining by 11 pp to 27 pp (with pp = percentage points) in terms of *MRR* on both datasets. Furthermore, *LAFF* is efficient; it takes at most 317 ms to provide suggestions for input instances of the proprietary dataset.

To summarize, the main contributions of this paper are:

- The LAFF approach, which addresses the problem of automated form filling for categorical fields, an important user interface challenge in many software systems. To the best of our knowledge, LAFF is the first work to combine BNs with local modeling and a heuristic-based endorser to provide accurate form filling suggestions, even for arbitrary filling orders and partially filled forms.
- An extensive evaluation assessing the effectiveness and efficiency of LAFF and comparing it with the state of the art.

The rest of the paper is organized as follows. Section 2 provides a motivating example and explains the basic definitions of automated form filling and its challenges. Section 3 introduces the basic machine learning algorithms used in this paper. Section 4 describes the different steps and the core algorithms of LAFF. Section 5 reports on the evaluation of LAFF. Section 6 discusses the usefulness, practical implications, and limitations of LAFF. Section 7 surveys related work. Section 8 concludes the paper.

2 DATA ENTRY FORM FILLING

In this section, we introduce the concepts related to data entry forms, provide a motivating example, define the problem of automated form filling, and discuss its challenges.

2.1 Data Entry Forms

A data entry form is typically composed of many fields (also called input parameters [77] or elements [2]), which can be of different types: textual, categorical, numerical, and file. Textual and numerical fields collect free text and numerical values respectively; users can freely input any value that is compliant with the field input validation rules. Categorical fields provide a list of candidate values from which the user has to choose (e.g., country); the source of candidate values is defined statically [1]. File fields are used to upload files, such as images and videos. During software design, these fields are associated with specific UI widgets, based on the corresponding type. For example, developers can use a list box or combo box to collect categorical data [75].

2.2 Motivating Examples

Although software users frequently fill forms, such activity is time-consuming and error-prone [64]. In the following, we describe two examples illustrating the main challenges faced while filling data entry forms with categorical fields.

Example 1: *Users require focused attention to choose options from categorical fields.*

Alison is a student majoring in biology. She uses information management platforms (such as NCBI [12] and SAIKS [28]) to record the basic information of biological samples. Given a biological sample (e.g., the genus *Pratylenchus*), she needs to fill fields “sex”, “tail shape”, and “species name” of the sample in a data entry form from such a system. All three fields are categorical with predefined values. After filling the first two fields, Alison starts to select the species name. However, the genus *Pratylenchus* currently includes over 80 valid species [28]. She has to scroll down the list to check for the relevant species. Although she can search a species by inputting the first letter of the name, many similar species names are presented (e.g., *pratensisobrinus*, *pseudocoffeae*, *pseudofallax*, and *pseudopratenensis*). She still requires focused attention to choose among these textually similar options in a limited time. According to existing studies, about half of the data entry errors are caused by selection error [61, 78]

Example 2: *Users require cognitive effort to match options with the actual value they plan to fill in.* The second example is inspired by a use case of our industrial partner. The use case refers to the

opening of a bank account for business customers; a simplified data entry form for this activity is shown on the left-hand side of Figure 1. We will use this data entry form as a running example to explain the form filling problem and illustrate our solution. For simplicity, let us assume that the data entry form contains only two fields: “company type” (textual) and “(company) primary field of activity” (categorical). When a company called *SmartLease* requests to open a bank account, the bank clerk Bill (who is the end-user interfacing with the data entry form) asks the *SmartLease* representative about the company type and inputs “leasing (company)” in the corresponding field; then, Bill selects “other financial services” for the “primary field of activity” field. Several weeks after the account opening, the data quality division of the bank detects a potential data quality issue regarding *SmartLease*, in the form of a mismatch between the actual activity implied by the company’s type and operations, and the company activity recorded when the account was open. This issue can be quickly solved by checking with *SmartLease* and amending the “primary field of activity” field with the correct value: “leasing services”. Nevertheless, such an issue can cause a business loss: for example, by knowing the actual “primary field of activity” of the company, the bank could have offered targeted products to its customer since the beginning of the business relation. Further investigation reveals that this issue occurred because, as a new employee in the bank, Bill was not familiar with all the 75 options defined in the “primary field of activity” field. He browsed the list for a limited time, compared the candidate values with the actual value he intended to fill in and finally selected an inappropriate value.

The aforementioned problems cannot be solved satisfactorily by existing solutions that support filling categorical fields, including those based on the search-by-keyword functionality, web browser plugins for autofill, as well as approaches that build field ontologies.

First, some solutions help users locate the candidate values in categorical fields with the search-by-keyword functionality. This functionality cannot solve the issues in the two examples, since users need to carefully compare textually similar values (as shown in Example 1) and have the burden to remember all the options to avoid searching an inappropriate value (as shown in Example 2).

Second, web browser plugins such as Chrome Autofill Forms [32] provide automatic form filling, but they simply reuse the inputs provided in past forms to automatically fill out fields in different forms with the same information (for example “zip code”). They do not leverage the knowledge provided by already filled fields to provide “intelligent” suggestions. Moreover, these tools are usually personalized for a single user, and cannot be used in the context of enterprise software systems, in which the end-user fills the same form with different information. As shown in Example 2, a bank clerk works daily with several bank accounts of different customers and cannot directly reuse the input instance of a customer to pre-fill an input form for another customer.

Third, some approaches automatically build ontologies for form filling [6, 9, 79]. They map a ‘target’ field (e.g., “zip code”) in a form to ‘source’ fields (e.g., “postal code” or “postcode”) in other filled forms to support data exchange across software systems. However, for domain-specific software systems (e.g., biology information management platforms), many fields are domain-specific (such as “tail shape” and “species name” in Example 1) and cannot be easily mapped to fields in other forms. In addition, due to legal or security policies, software systems for governments and enterprises have constraints on sharing records across systems (as the banking system in Example 2).

For all the above reasons, there is a need to design a semi-automated method that developers can adopt to support and guide users during the form filling activity.

2.3 Problem Definition

In this paper, we deal with the automated form filling problem, which can be informally defined as the problem of suggesting possible values for the form fields a user is about to fill in, based on

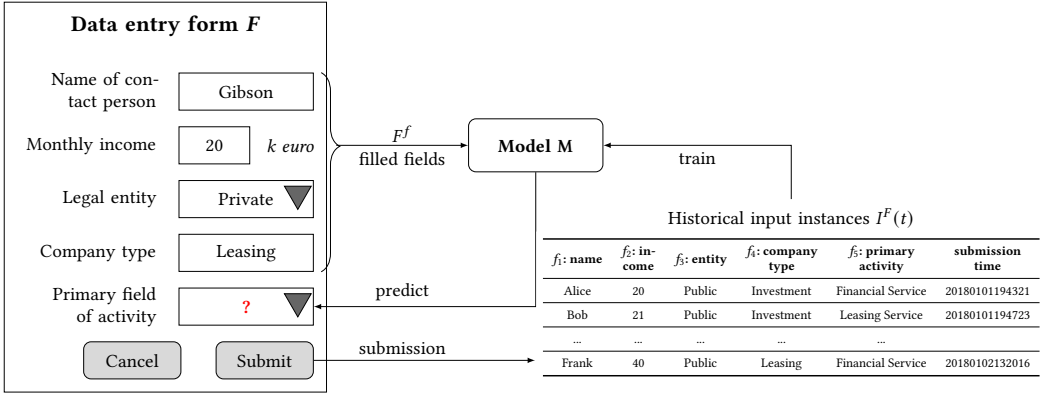


Fig. 1. The Automated Form Filling Problem

the values of the other fields and on the values provided as input in previous data entry sessions. We target *categorical fields* for automated filling since they require cognitive effort and focused attention for users to choose among the (typically) large set of options. The task of filling categorical fields may be slow and frustrating to users [33], and may lead to data quality issues, as shown in the examples above. We define the automated form filling problem as follows.

Let F be a data entry form with n fields $F = \{f_1, f_2, \dots, f_n\}$; each field f_i can take values from a domain V_i (which always includes a special element \perp representing an empty field); let $F^c \subseteq F$ be the set of categorical fields.

When a form F is being filled, at any time the fields can be partitioned in two groups: fields that have been already filled (denoted by F^f) and unfilled fields (denoted by F^u); we have that $F^f \cup F^u = F$ and $F^f \cap F^u = \emptyset$.

When a filled form F is about to be submitted (e.g., to be stored in a database), we define an *input instance* of F $I^F = \{\langle f_1, v_1 \rangle, \dots, \langle f_n, v_n \rangle\}$ with $f_i \in F$ and $v_i \in V_i$, as the set of pairs $\langle \text{field}, \text{value} \rangle$ from F ; we use the subscript t_j as in $I_{t_j}^F$ to denote that the input instance I^F was submitted at time t_j . We use the notation $I^F(t)$ to represent the set of *historical input instances* of form F that have been submitted up to a certain time instant t ; $I^F(t) = \{I_{t_i}^F, I_{t_j}^F, \dots, I_{t_k}^F\}$, where $t_i < t_j < t_k < t$. Hereafter, we drop the superscript F when it is clear from the context.

The automated form filling problem can be defined as follows. Given a (partially filled) form $F = F^f \cup F^u$, a set of historical input instances $I^F(t)$, and a target field $f_p \in (F^u \cap F^c)$ to fill, we want to build a model M that at time t can predict a value v_p for f_p based on F^f and $I^F(t)$. Notice that in this problem definition the filling order of the fields in F is unrestricted.

Application to the running example. Figure 1 shows an example explaining the automated form filling problem. We have a data entry form F for a banking system with five fields: f_1 : "name of contact person", f_2 : "monthly income", f_3 : "legal entity", f_4 : "company type", and f_5 : "primary field of activity". Among them, fields "legal entity" and "primary field of activity" are categorical (i.e., $F^c = \{f_3, f_5\}$). During the data entry session, users provide values for these fields, which are stored in a database upon submission of the form. The table on the right-hand side of Figure 1 shows some historical input instances filled by the bank customers through the data entry form; the submission timestamp t of these input instances is automatically recorded. In the table, each row represents an input instance (e.g., $I_{20180101194321}^F = \{\langle \text{"name"}, \text{Alice} \rangle, \dots, \langle \text{"primary activity"}, \text{Financial Service} \rangle\}$);

the column names correspond to the field names in the data entry form¹. With these historical input instances, we can build a model M to learn the relationships of values filled in fields f_1 to f_5 by different customers. Notice the submission timestamp is not used for model building; it is only used for distinguishing different input instances.

Continuing the example, let us assume that, at a certain point of the data entry session, a customer has provided the values *Gibson*, *20*, *Private*, and *Leasing* for fields f_1 to f_4 respectively, as shown on the left-hand side of Figure 1; the unfilled field f_5 (“primary field of activity”) is the next (categorical) field to fill in. Our goal is to use the model M to predict the value of f_5 based on the values of the filled fields f_1 to f_4 .

2.4 Challenges of Automated Form Filling

Several automated form filling approaches have been proposed [38, 53, 71]; the basic idea is to mine dependencies among fields from the values recorded in previous form filling sessions, to build recommendation models. These models can then be used to suggest possible values on a target field based on the filled fields in the current form. Nevertheless, state-of-the-art approaches exhibit some limitations when dealing with (1) forms filled following an arbitrary order, and (2) partially filled forms.

First, while filling in a data entry form, it is very frequent to have little or no restriction on the order of user’s inputs. A user may select any field as the next target or even go back and modify already filled fields. In other words, the set of filled fields (F^f) and the target field to suggest ($f_p \in F^u$) keep changing. This scenario is different from the one considered by many recommender systems in the software engineering domain [54, 55], in which models are trained on predefined features/attributes (e.g., code metrics) to predict a specific target (e.g., source code defects). Some automated form filling approaches [38] require a fixed form filling order before building the recommendation models. However, this assumption is unrealistic from a practical standpoint. Although some approaches [53, 71] are insensitive to the form filling order (e.g., suggesting possible values of a target field based on the frequency of values in historical inputs), they may not provide accurate suggestions due to their limitations in accurately mining dependencies among fields (as discussed in section 5.2). Hence, one of the challenges in automated form filling is how to build recommendation models (e.g., by mining dependencies among fields) without making any assumption on the order in which fields are filled.

Second, at a certain time during the data entry session, a form is usually partially filled: this means that a recommender system for automated form filling can only use the knowledge in currently filled fields (F^f). However, when the filled fields do not provide enough knowledge to predict the target—based on our preliminary experiments—existing approaches [38, 53] yield inaccurate suggestions. The challenge, when dealing with partially filled forms, is how to discard low-confidence suggestions, in order to make suggestions only when a high degree of confidence is achieved.

3 PRELIMINARIES

Before illustrating our approach, we first briefly introduce two basic machine learning algorithms we rely on.

¹In the case of this illustrative example, we assume that the mapping between field names and column names can be retrieved in some way, for example by manually analyzing the existing software design documentation or software implementation. We provide more explanations on identifying such a mapping in section 6.2.1.

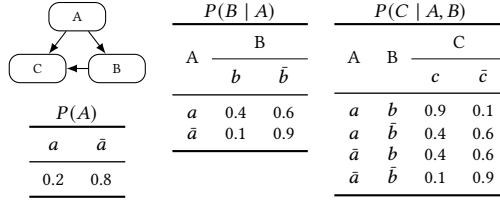


Fig. 2. An Example of BN and the Probability Functions of its Nodes

3.1 Bayesian Networks

Bayesian networks (BNs) are probabilistic graphical models (PGM) in which a set of random variables and their conditional dependencies are encoded as a directed acyclic graph: nodes corresponds to random variables and edges correspond to conditional probabilities.

The use of BNs for supervised learning [27] typically consists of two phases: structure learning and variable inference.

During *structure learning*, the graphical structure of the BN is automatically learned from a training set. First, the conditional probability between any two random variables is computed. Based on these probabilities, optimization-based search (e.g., hill climbing [29]) is applied to search the graphical structure. The search algorithm initializes a random structure, and then iteratively adds/deletes its nodes and edges to generate new structures. For each new structure, the search algorithm calculates a fitness function (e.g., Bayesian information criterion, BIC [62]) based on the nodes' conditional probabilities and on Bayes' theorem [27]. Structure learning stops when it finds a graphical structure that minimizes the fitness function.

Figure 2 shows an example of the BN structure with three random variables: variable B depends on variable A ; variable C depends on variables A and B . In the PGM, each node is associated with a probability function (in this case encoded as a table), which represents the conditional probability between the node and its parent(s). For example, in Figure 2 each variable has two values; the probability table for B reflects the conditional probability $P(B | A)$ between A and B on these values.

Variable inference infers unobserved variables from the observed variables and the graphical structure of the BN using Bayes' theorem [27]. For example, we can infer the probability of $C = c$ when the value of A is a (denoted as $P(c | a)$) as follows:

$$\begin{aligned}
 P(c | a) &= \frac{P(a, c)}{P(a)} = \frac{P(a, b, c) + P(a, \bar{b}, c)}{P(a)} \\
 &= \frac{P(c | a, b)P(b | a)P(a) + P(c | a, \bar{b})P(\bar{b} | a)P(a)}{P(a)} \\
 &= \frac{0.9 * 0.4 * 0.2 + 0.4 * 0.6 * 0.2}{0.2} = 0.6
 \end{aligned}$$

BNs have been initially proposed for learning dependencies among discrete random variables. They are also robust when dealing with missing observed variables; more specifically, variable inference can be conducted when some conditionally independent observed variables are missing [27].

3.2 K-modes

K -modes is a clustering algorithm that extends the k -means one to enable clustering of categorical data [41]. The algorithm first randomly selects k instances in the data set as the initial centroids.

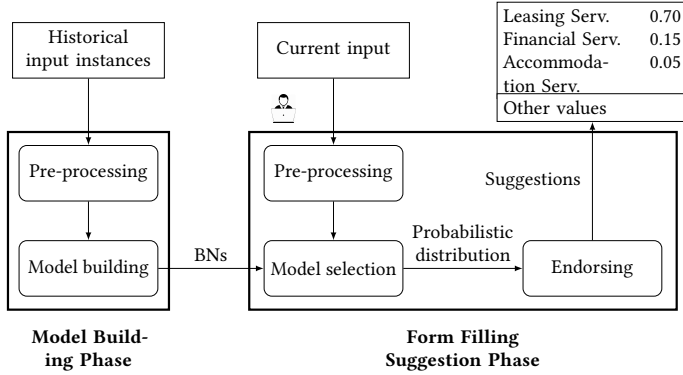


Fig. 3. Main Steps of the LAFF Approach

Each instance is represented with a vector of categorical attributes. The algorithm clusters the instances in the dataset by calculating the distances between the instances and each centroid. The distance, also called dissimilarity measure, is defined as the total mismatches of the corresponding categorical attributes of an instance and of a centroid. Based on the clustering results, new centroids are selected, which represent the modes of categorical attributes in each cluster. The algorithm then re-clusters the instances according to the new centroids. This process is repeated until the centroids remain unchanged or until it reaches a certain number of iterations.

4 APPROACH

In this section, we present our machine-learning based approach for form filling, named LAFF (Learning-based Automated Form Filling).

LAFF includes two phases: *model building* and *form filling suggestion*, whose main steps are shown in Figure 3. In the former, LAFF analyzes historical input instances of a data entry form and uses dependency analysis to build BNs that represent the fields in the form and their conditional dependencies. This phase occurs offline, before deploying LAFF as an assistive technology for data entry. The *form filling suggestion* phase occurs during the data entry session: given a target field, LAFF selects a BN among those built in the *model building* phase and predicts possible values based on the values in the already-filled fields and their conditional dependencies; the predicted values and the corresponding predicted probability (endorsed based on field dependencies and prediction confidence) are then provided to the end-user as suggestions.

4.1 Pre-processing

Both phases of LAFF include a pre-processing step to improve the quality of the form filling data; this step is based on best practices for predictive data mining [3].

Typically, historical input instances have many *missing values* due to the presence of optional fields in input forms. Fields that have a high number of missing values do not provide representative information for model building; hence, they can be removed. We remove fields for which $T_p^v\%$ or more of the values are missing, where T_p^v is a user-configurable threshold (with default value equal to 90).

We also remove file and textual fields that have a high number of unique values, since they typically correspond to form fields for which users frequently provide new string values (e.g., the textual field “name”). To identify such fields, we compute the ratio of unique values of a field in the

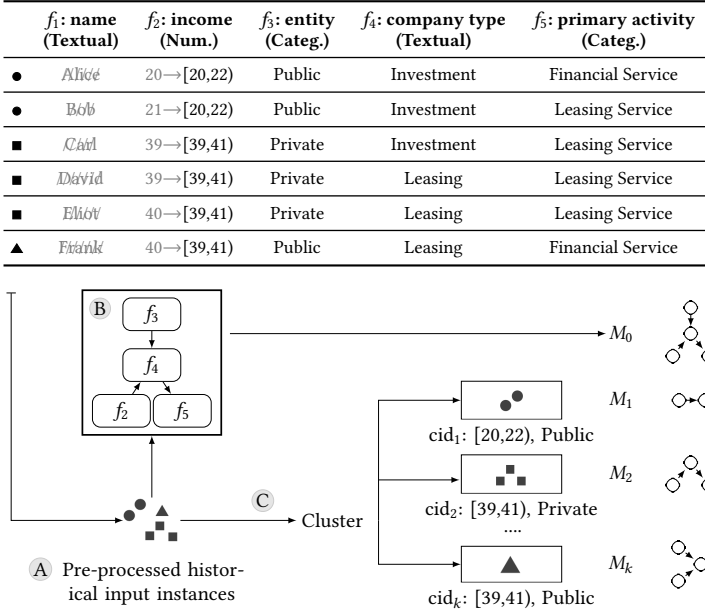


Fig. 4. Example of Model Building on Pre-processed Historical Input Instances

historical input instances; if the ratio is larger than a user-configurable threshold T_p^u (with default value equal to 0.9), the corresponding field is removed.

Furthermore, we delete a historical input instance if more than $T_p^m\%$ of its field values are missing, where T_p^m is a user-configurable threshold (with default value equal to 50). After deletion, we perform data imputation [44] on the remaining data that exhibit missing values. Numerical fields are imputed using the mean value of this field; categorical and textual fields are imputed using a default label “UNKNOWN”.

We also apply data discretization to numerical fields to reduce the number of unique values. Numerical values are transformed into discrete intervals based on information gain analysis, a widely used discretization method first proposed in decision trees [13].

During the data entry session, we ignore values in the fields that were removed in historical input instances, and map numbers onto intervals.

Application to the running example. The table at the top of Figure 4 shows an example of historical input instances filled through the data entry form in Figure 1. Each row is a historical input instance filled by a user. During pre-processing, LAFF removes the field “name” since all its values are unique (we crossed out the text of these values with a hatch pattern to represent the removal). Also, the values of field “income” are discretized into intervals. During the data entry session, as shown in Figure 1, a user fills the fields “name” with *Gibson*, “income” with 20, “legal entity” with *Private*, and “company type” with “Leasing”; “primary activity” is the next field to be filled. Through the application of the pre-processing steps, LAFF ignores the value for field “name” and maps the value 20 of field “income” to the interval [20, 22).

Algorithm 1: Model Building

Input: Pre-processed historical input instances $I^F(t)'$
Output: List of probabilistic graphical models \mathcal{M}
Historical input instance clusters C

```

1  $\mathcal{M} \leftarrow$  empty list;
2  $M_0 \leftarrow \text{trainBayesianNetwork}(I^F(t)')$ ;
3  $\mathcal{M}.\text{append}(M_0)$ ;
4 independent field set  $F^I \leftarrow \emptyset$ ;
5 foreach field  $f_i \in M_0$  do
6   if  $\text{getParents}(M_0, f_i) = \emptyset$  then
7      $F^I \leftarrow \{f_i\} \cup F^I$ ;
8   end
9 end
10 number of cluster  $k \leftarrow \text{elbowMethod}(I^F(t)', F^I)$ ;
11  $C = \{I^F(t)'_1, \dots, I^F(t)'_k\} \leftarrow \text{kModes}(I^F(t)', F^I, k)$ ;
12 for  $i \leftarrow 1$  to  $k$  do
13    $M_i \leftarrow \text{trainBayesianNetwork}(I^F(t)'_i)$ ;
14    $\mathcal{M}.\text{append}(M_i)$ ;
15 end
16 return  $\mathcal{M}, C$ ;
```

4.2 Model building

The goal of the *model building* phase is to mine dependencies from historical input instances of a data entry form.

Due to the arbitrary order for filling the form, the filled fields and the target field keep changing. When we take the filled fields as features to predict the target field, the arbitrary form filling order results into a large set of feature-target combinations. For example, let us consider a data entry form with n fields, with $t \leq n$ of them being categorical and thus representing the possible targets. When we take one of the categorical fields as the target, assuming that a random order is used for form filling, in principle users may fill any subset of the remaining $n - 1$ fields, resulting in a total of $2^{n-1} - 1$ possible combinations of filled fields (i.e., features). The total number of feature-target combinations is equal to $t * (2^{n-1} - 1)$. Normally, a model would need to be trained on each target-features combination to ensure the assumption of identical features and target [21] in the model building and form filling suggestion phases. As we will show through our evaluation in section 5, adopting such an approach would require to train more than 220 000 models on one of our datasets. The total time required to train this large number of models would be impractical for a production-grade system.

To solve this problem, we capture dependencies with BNs, in which variables correspond to form fields. By using BNs, we can analyze the dependency between filled fields and target fields without training models on specific combinations of features (i.e., filled fields) and target field. In addition, as mentioned in section 3.1, BNs are robust when dealing with missing values. This means that even when a data entry form is partially filled, BNs can still infer the probability distribution of target fields using only the information in the filled fields and the underlying PGM.

In this work, we learn the structure of BN from the pre-processed historical input instances. Following the workflow of BN presented in section 3.1, we represent each field in the historical input instances as a random variable. BN computes the conditional probability between any two fields and uses a search-based optimizer to automatically optimize the structure of BN based on the conditional probability of fields and the fitness function. In this study, we use hill climbing as the

optimizer, because it shows a good trade-off between computational demands and the quality of the models learned [29]. We define the fitness function in terms of BIC [62], which aims at best fitting the data, while avoiding over-fitting by complex structures. The element denoted with **B** in Figure 4 shows an example of the BN structure learned from the data in block **A** (where the different black shapes correspond to the various rows in the table at the top of Figure 4).

Algorithm 1 illustrates the main steps of this phase. LAFF takes as input the pre-processed historical input instances $I^F(t)'$ as the training data to mine field dependencies. Initially, we learn the BN over the entire training data (line 2). This global model, denoted as M_0 , represents the general dependencies among fields. However, historical input instances may form different groups that share similar characteristics. For example, in the historical input instances of Figure 4, users having the same value for fields “income” and “legal entity” may share specific values for fields “company type” and “primary activity”. The global model M_0 may not learn the fine-grained field dependencies for specific values of “income” and “legal entity” due to the influence of input instances with other values for those fields. For example, using the entire dataset in Figure 4, one could determine that the conditional probability of having “primary activity” equal to *Leasing Service* when “Company type” is *Leasing* is 66.7%. However, this conditional probability increases to 100%, if we only consider the input instances where “income” is in the range [39, 41) and “legal entity” is equal to *Private*. Hence, LAFF trains local models on subsets of $I^F(t)'$ to learn fine-grained dependencies.

More specifically, LAFF first selects a set of fields F^I that are independent from other fields in the probabilistic graph of M_0 (lines 4–9). For example, in block **B** of Figure 4, fields f_2 and f_3 are selected as they do not depend on other parent nodes (fields). We use the fields in F^I as the main fields to form partitions of $I^F(t)'$ having similar characteristics for two reasons. First, these fields are not intercorrelated since they do not directly and strongly depend on each other. Second, these fields are root nodes and influence the values of other fields in the BN; when the values on these fields are similar, we are likely to obtain a similar probability distribution for the other fields.

LAFF produces (lines 10–11) a set C of clusters of $I^F(t)'$ based on the fields in F^I . We assume that the clustering process reduces the data variation of $I^F(t)'$: models trained on these data, which show less statistical variation, may provide more accurate suggestions even when the size of each cluster is smaller than that of $I^F(t)'$. This process is called local modeling; it has been already applied in software engineering, e.g., to cluster software projects and mine project-specific relationships of software metrics [55].

To extract clusters from $I^F(t)'$, LAFF represents each historical input instance as a tuple of the form $\langle \text{values in } F^I, \text{input instance} \rangle$. It clusters these tuples based on the values in F^I using the k -modes algorithm. The optimal number of clusters k is automatically determined with the elbow method. LAFF runs k -modes within a range of k values (e.g., from 1 to 100) and determines the value of k that minimizes the average within-cluster distance with the cluster centroids (denoted with “cid” in block **C** of Figure 4). After clustering, LAFF trains a local BN model M_i (lines 12–15) based on the input instances in each cluster. These *local* models, denoted M_1, \dots, M_k , capture specialized field dependencies on partitions of $I^F(t)'$.

The algorithm ends by producing a list \mathcal{M} of BNs, where $\mathcal{M} = [M_0, M_1, \dots, M_k]$, and the set C of clusters of the historical input instances.

Application to the running example. Initially, LAFF trains a global model M_0 with the historical input instances in block **A** in Figure 4. Block **B** of Figure 4 shows an example of the learned field dependencies. Based on M_0 , LAFF selects fields f_2 : “income” and f_3 : “legal entity” as the main fields for local modeling since they do not depend on other parent nodes (fields). LAFF clusters the historical input instances according to fields “income” and “legal entity” (block **C** of Figure 4).

Algorithm 2: Form Filling Suggestion

Input: Models $\mathcal{M} = [M_0, M_1, \dots, M_k]$
 Clusters $C = \{I^F(t)_1', I^F(t)_2', \dots, I^F(t)_k'\}$
 Filled fields $F^f = \{(f_1^f, v_1^f), \dots, (f_m^f, v_m^f)\}$
 Target field f_p
 Number of suggested values n_r
 Threshold θ

Output: List of predicted values V_p for f_p

```

1  $F^{f'} \leftarrow \text{getPreprocessedData}(F^f)$ ;
2  $D = \{d_1, \dots, d_k\} \leftarrow \text{calcClusterDistance}(C, F^{f'})$ ;
3 Model  $M_{cur} \leftarrow \mathcal{M}.\text{get}(M_0)$ ;
4 if  $\text{getNumOfMinDistance}(D) = 1$  then
5    $i \leftarrow \text{getMinDistanceID}(D)$ ;
6    $M_{cur} \leftarrow \mathcal{M}.\text{get}(M_i)$ ;
7 end
8 List of Pairs  $\langle v_p, pr \rangle$  of candidate values and probability distribution
    $Candidates = \text{predictCandidates}(M_{cur}, F^{f'}, f_p)$ ;
9  $Candidates^R \leftarrow \text{getTopRanked}(Candidates, n_r)$ ;
10 Bool  $checkDep \leftarrow \text{isMember}(\text{getParents}(M_{cur}, f_p), F^f)$ ;
11 Bool  $checkProb \leftarrow (\text{getSumProb}(Candidates^R) > \theta)$ ;
12 if  $checkProb \vee checkDep$  then
13   foreach  $v_{p_i}$  s. t.  $\langle v_{p_i}, pr_i \rangle \in Candidates^R$  do
14      $V_p.append(v_{p_i})$ ;
15   end
16 end
17 return  $V_p$ ;

```

Three clusters are automatically identified with centroids “[20, 22), Public”, “[39, 41), Private”, and “[39, 41), Public” ($k=3$). We use circular, rectangular, and triangular icons to represent the historical input instances belonging to different clusters. LAFF trains three local models M_1 , M_2 , and M_3 , based on these clusters; these three models are three distinct BNs capturing specialized field dependencies (as shown on the right of Figure 4). After the model building phase, LAFF outputs four models: a global BN model M_0 and the three local BN models M_1 , M_2 , and M_3 .

4.3 Form Filling Suggestion

The *form filling suggestion* phase occurs during the data entry session and assumes that the models in \mathcal{M} , built in the *model building* phase, are available. Given a target field f_p , LAFF selects a BN model $M \in \mathcal{M}$ and predicts possible values of f_p based on the already-filled fields F^f and their conditional dependencies captured in M . The main steps of the *form filling suggestion* phase are shown in Algorithm 2.

The algorithm takes as input a list of models \mathcal{M} , a set of clusters C , a set F^f of already-filled fields with their values, a target field f_p , and some auxiliary parameters representing the number of expected suggestions for f_p and an endorsing threshold. After pre-processing the filled fields in F^f using the techniques discussed in § 4.1 and obtaining the new set $F^{f'}$, LAFF computes the distance between the filled fields $F^{f'}$ and each cluster in C (line 2). The distance is defined as the dissimilarity measure adopted in the k -modes clustering algorithm used in the model building phase; it is the total number of mismatches between $F^{f'}$ and the centroid of each cluster on the corresponding fields.

LAFF attempts to select a local model from M_1, \dots, M_k , corresponding to the cluster with minimal distance, to predict the target field, since this model may capture the fine-grained characteristics of F^f . In our example, this could be a model trained on the instances that have the same values for the fields “legal entity” and “income” in F^f . However, a unique and optimal local model cannot always be found: given a partially filled data entry form, there could be cases for which the distance of the filled fields to different centroids is equal. For example, in Figure 4, local models M_2 and M_k are specialized for different values of the field “legal entity” (i.e., “Private” and “Public”) but assume the same value for the field “income” ([39, 41]). Let us consider the case in which the set F^f contains only the field “income” (with a value equal to “40”) and the target field is “primary activity”. In this case, we cannot reliably select between the two models M_2 and M_k for prediction, since we have insufficient information to decide which model is “more local” (i.e., specialized) for this input (i.e., the distance to the two centroids is the same). One possible solution for this problem is ensemble learning, which considers the predictions of both models jointly (e.g., bagging) [83]. However, this solution could significantly increase prediction time. Specifically, depending on the deployment configuration, in the worst case, the ensemble prediction time would be the sum of the prediction time of all k local models (i.e., all the local models are not specialized for the current input), which may exceed the acceptable response time for a practical application, as presented in section 5.3 and further discussed in section 6.2.1. Given the interactive nature of data-entry applications, having a short prediction time is important. Hence, when LAFF finds more than one minimal distance and no single cluster is particularly suited for the current input, it selects M_0 for prediction, since it is trained with the entire set of historical input instances (lines 3–7).

After selecting the most appropriate model for prediction, LAFF predicts the candidate values for the target field (line 8) and ranks the topmost n_r values, according to their probability distribution (line 9).

Endorsing. During the data entry session, the filled fields in the current input instance do not always provide enough information to predict values for the target field, leading to inaccurate suggestions. Such a situation can occur because of two reasons. One reason is that the filled fields may not have explicit dependencies with the target field, according to the probabilistic graph. For example, in Figure 4, f_2 is independent from f_3 ; LAFF will not accurately infer f_3 merely with the knowledge of f_2 . Another reason is that there may not be enough historical input instances to learn the conditional probability between two fields for specific values. For example, in the example in Figure 4, we have no historical input instance with values of field “income” greater than 41; such value provides limited information to infer other fields.

In the context of automated form filling, users might be reluctant to use an automated form filling tool, if the tool provides many inaccurate suggestions which users can hardly find the correct value they intend to fill in. To avoid such a situation, LAFF includes a *heuristic-based endorser*, which decides whether the suggestions determined in the previous step are accurate enough to be returned to the user.

To deal with the first cause of inaccurate suggestions, LAFF analyzes the dependency between the filled fields and the target field (line 10), to check whether the target field directly depends on one of the filled fields in the BN. More precisely, function *getParent* computes a list of parent fields the target field directly depends on; function *isMember* checks whether any of the filled fields is in the parent field list. The result of this check is saved in the Boolean flag *checkDep*, which is true when the target field directly depends on one of the filled fields. A direct dependency indicates that the filled fields can reliably determine the value of the target field.

To deal with the second cause of inaccurate suggestions, LAFF analyzes the predicted probability distribution of the values for the target field. For a probability model like BN, the probability of

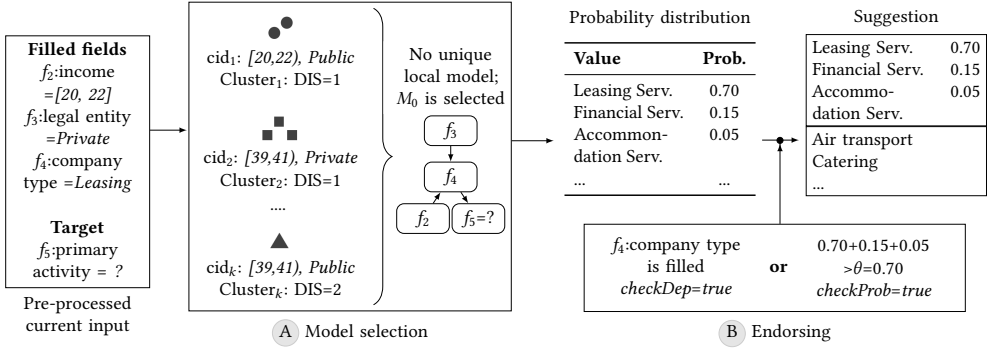


Fig. 5. Workflow for Form Filling Suggestion

each value is inferred based on the information from the filled fields. LAFF computes the sum of the top- n_r probability values in the distribution through function *getSumProb*. If this value is larger than a user-defined threshold θ , it means LAFF may have enough information for variable inference; the result of this check is saved in the Boolean flag *checkProb* (line 11). From a practical standpoint, threshold θ reflects how much uncertainty users are willing to accept regarding the suggestions provided by LAFF.

If one of the flags *checkProb* and *checkDep* evaluates to *true*, LAFF populates the list of suggested values to be returned to the user based on the top-ranked candidate values; otherwise, LAFF returns an empty list (lines 12–15). For example, assuming $\theta = 0.70$, $n_r = 3$, and the probability for the top-3 candidate values as shown in the top right corner of Figure 3, the sum of the top-3 probability values (returned by *getSumProb*) is $0.70 + 0.15 + 0.05 = 0.90$; *checkProb* corresponds to the evaluation of $0.90 > 0.70$, which is *true*; hence, LAFF decides to yield the list of suggestions to the user.

Application to the running example. Given the new input instance shown on the left side of Figure 5 (i.e., the instance “income”=[20, 22], “legal entity”=Private, and “company type”=Leasing, as obtained after pre-processing), LAFF suggests the possible values of “primary activity”. As shown in block A of Figure 5, LAFF first attempts to select a unique local model by calculating the distance between the current input instance and the centroids of the three clusters generated in block C of Figure 4; however, such a local model cannot be found because the distances with “ cid_1 ” and “ cid_2 ” are both 1. Hence, LAFF uses M_0 for prediction. According to the variable inference method in BN (explained in section 3.1), LAFF outputs the probability distribution of the candidate values for the field “primary activity”. Let us assume, as an example, that the probability distribution is “Leasing Service=0.70, Financial Service=0.15, Accommodation Service=0.05, ...”. By means of the endorser module (block B of Figure 5), LAFF uses this probability distribution to decide whether to present the suggestions to the user. For example, let us further assume the data quality engineers in the bank set θ to 0.70 and configures LAFF to suggest three values. On the one hand, LAFF finds that the target field f_5 : “primary activity” directly depends on f_4 : “company type”, which was already filled by the user; the *checkDep* flag is *true*. On the other hand, the sum of the top-3 probability values is 0.90, which is higher than θ ; the *checkProb* flag is *true*. Since the endorser module endorses a suggestion when one of these two flags evaluates to true, LAFF provides a suggestion to the user: the three values above are put to the top of the list while the other candidate values are presented in their original order (e.g., alphabetically).

Table 1. Information about the Fields in the Datasets

Dataset	# of fields	# of instances	Name of categorical fields (# of candidate values)	Value frequency	
				top-1	top-5%
NCBI	26	74105	sex(3), tissue(68), cell-line(50), cell-type(63), disease(84), ethnicity(40)	40.8%	59.4%
PROP	33	174446	title(18), sex(3), legal capacity(7), country(208), first nationality(206), civil status(8), matrimonial regime(6), activity(13), status(15), function(41), contract(8), field of activity(75), primary activity(3), country of activity(198)	48.4%	65.6%

5 EVALUATION

In this section, we report on the evaluation of our approach (LAFF) for automated form filling. First, we assess the overall accuracy of LAFF in suggesting appropriate values to automatically fill in the fields of data entry forms, and compare it with state-of-the-art form filling algorithms. We also assess the performance of LAFF, in terms of training time and prediction time, for practical applications. Then, we evaluate how the use of local modeling (in the *model building* phase) and heuristic-based endorser (in the *form filling suggestion* phase) affect the accuracy of LAFF. Last, we assess the impact of the number of filled fields and the size of the training set on the effectiveness of LAFF.

More specifically, we evaluated LAFF by answering the following research questions:

- RQ1 *Can LAFF provide accurate suggestions for automated form filling, and how does it compare with state-of-the-art algorithms?*
- RQ2 *Is the performance of LAFF (in terms of training time and prediction time) suitable for practical application in data-entry scenarios?*
- RQ3 *What is the impact of using local modeling and heuristic-based endorser on the effectiveness of LAFF?*
- RQ4 *What is the impact of the number of filled fields on the effectiveness of LAFF?*
- RQ5 *What is the impact of the size of the training set on the effectiveness of LAFF?*

5.1 Dataset and Settings

Datasets. We evaluated LAFF using a public dataset in the biomedical domain (dubbed NCBI) and a proprietary dataset, extracted from a production-grade enterprise information system, provided by our industrial partner (dubbed PROP).

The NCBI dataset contains the metadata for diverse types of biological samples from multiple species [12]. We selected this dataset because it has been used in a previous study on metadata suggestion for biomedical datasets [53], which provided also the design of the corresponding data entry form in the CEDAR workbench [31]. More specifically, following the evaluation methodology described in [53], we considered the subset of the NCBI dataset related to the species “Homo sapiens” and the corresponding data entry form based on the specification of the BioSample Human package v1.0². We downloaded the dataset from the official NCBI website³. In the dataset, the data is organized as a table. Each row is an input instance filled by a user. Retrieving the mapping between column names and field names was trivial since the column names in the dataset are the same as the field names. As shown in Table 1, the NCBI dataset has 26 fields, six of which are categorical. These categorical fields have between 3 and 84 candidate values to be selected by users. We calculated the frequency by which users select different values during form filling: the most

² <https://submit.ncbi.nlm.nih.gov/biosample/template/?package=Human.1.0&action=definition>

³ <https://ftp.ncbi.nlm.nih.gov/biosample/>

frequent (i.e., top-1) and the top-5% most frequent values are selected, on average, respectively in 40.8% and 59.4% of the instances for different categorical fields. Given the sparseness of the dataset (caused by the optional fields), as suggested in [53], we identified the empty values (e.g., “n/a”, “null”), and only retained records with at least three fields (out of six) with non-empty values; in total, the NCBI dataset contains 74 105 input instances.

The PROP dataset contains customer data that are provided through a web-based data entry form, which is filled out upon creation of a new customer account. We extracted the dataset from the distributed database of our industrial partner, where all the input instances of a certain form are organized as a database table. Each row in the table is an input instance and each column represent a form field. We identified the mapping between the column names in the table and the field names in the data entry form by consulting the available software documentation. As shown in Table 1, the PROP dataset has 33 fields, 14 of which are categorical (with the number of candidates values ranging from 3 to 206). In terms of frequency according to which users select different candidate values, the top-1 and the top-5% most frequent values are selected, on average, respectively in 48.4% and 65.6% of the instances. According to the form design, eight of the categorical fields are mandatory to be filled; hence, we do not remove spare records as done for the other dataset; in total, the PROP dataset contains 174 446 input instances.

We remark that both datasets represent the input instances from real-world data entry forms (i.e., the NCBI platform and a production-grade enterprise information system). The number of fields in these systems is comparable with or larger than the data entry forms used in the related work. For example, we calculated the average number of fields of data entry forms in the TEL-8 dataset, a manually collected dataset with 447 web forms (with no input instances), which is used in the literature on form filling [9, 45]. In this dataset, each form has 6.39 fields on average. The data entry forms in our study are more complex, ranging from 26 to 33 fields of different types.

Dataset Preparation. For the two datasets, as discussed in section 2.3, all the categorical fields are the targets for automated form filling. However, we excluded the fields with less than 10 candidate values (e.g., “sex”, which has only three values in both datasets) as users may easily browse all the values in these fields, without the need for form-filling automation. The threshold for excluding fields was determined together with the data quality engineers and some data entry operators of our partner. We find the majority of categorical fields we evaluated are related to certain domains or business processes; they include fields “tissue”, “cell-line”, “cell-type”, “disease” and “ethnicity” for the biological domain, and fields “activity”, “status”, “function”, “field of activity”, and “country of activity” for the financial domain. These fields are more difficult to fill than basic user information (e.g., name, sex, and age), since users need to understand the meaning of candidate values.

Since both datasets automatically recorded the submission time of each input instance. we split the dataset into two subsets containing 80% and 20% of input instances based on their submission time, used respectively for training and testing. The input instances (excluding the information of the submission time) in the training set are used to train LAFF. As for the testing input instances, since there is no information on the actual filling order used to input the data, we considered two form filling orders to simulate the data entry session. More specifically, we simulated two types of filling scenarios: “sequential filling” and “random filling”. The former corresponds to filling data entry forms in the default order, as determined by the form *tab sequence*, e.g., the navigation order determined by the HTML attribute `tabindex` in web UI designs [26]. It simulates the logical order many users follow to fill out forms, especially when they use a keyboard to navigate form fields [56]. The latter represents the scenario when users may select any field as the next target, and even go back to modify already filled fields. These two form filling orders represent two

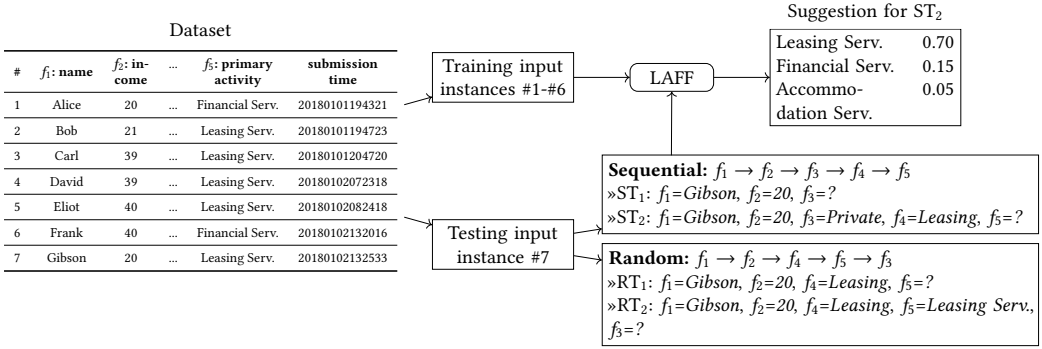


Fig. 6. Example of Dataset Preparation (Training and Testing Sets)

opposite extremes⁴ of user behavior during a real data-entry session. We simulated random filling by randomly generating an order for each testing input instance. In both form filling scenarios, the filled fields considered by LAFF are the fields that precede each target. For each target field, we consider the actual value filled by the user as the ground truth.

Dataset Preparation - Example of Application. Figure 6 shows an example of application of our dataset preparation method for the training and testing sets. Let us consider a dataset containing seven input instances submitted through a data entry form, shown on the left-hand side of Figure 6. Following the running example introduced in section 2.3, the form has five fields, two of which are categorical (e.g., f_3 : “legal entity” and f_5 : “primary activity”). We split the dataset into the training and testing sets according to the submission time: we take 80% of input instances (i.e., #1-#6) to train LAFF; the testing set contains the remaining 20% of the instances, in this case input instance #7. The testing set is further processed to simulate the two types of filling scenarios. When using the sequential filling order, users fill the data entry form following the tabindex of fields in the form (e.g., from f_1 to f_5 sequentially): starting from the input instance #7, we generate test instances ST₁ and ST₂ for categorical fields f_3 and f_5 , respectively. For each categorical field (i.e., the target), the actual value filled by the user is the ground truth (e.g., the ground truth for the field f_5 is ‘Leasing Serv.’). When using the random filling order, we randomly generate a field order for each input instance (e.g., $f_1 \rightarrow f_2 \rightarrow f_4 \rightarrow f_5 \rightarrow f_3$ for the input instance #7); based on this order, we then generate test instances RT₁ and RT₂.

Implementation and Settings. We implemented LAFF as a Python program; we used the open-source library pgmpy [8] for working with Bayesian networks.

We configured LAFF (through parameter n_r in Algorithm 2) to suggest the top 5%, most likely values for each target field. Based on the number of candidate values for each field in the datasets (indicated in parentheses in the rightmost column of Table 1), suggesting the top 5% values means showing between one (for field “activity” in the PROP dataset) and ten (for field “country” in the PROP dataset) suggested values to users. This is in accordance with other recommender systems in software engineering, in which only a list of few candidates is suggested for consideration [59, 80].

⁴For some large data entry forms, UI designers can semantically partition related fields into sections. Users can then move between sections in sequential order, while using the random order to fill fields within a section. This is a “middle-ground filling” order, which sits between “sequential filling” and “random filling”. We have not evaluated this scenario since it requires additional knowledge about the partitioned sections, which was not available for the two datasets we have considered.

We set the threshold θ to 0.7 based on the feedback received by data entry operators and data quality engineers of our partner.

We performed the experiments on the NCBI dataset with a computer running macOS 10.15.5 with a 2.30 GHz Intel Core i9 processor with 32 GB memory. As for the experiments on the PROP dataset⁵, we performed them on a server running CentOS 7.8 on a 2.60 GHz Intel Xeon E5-2690 processor with 125 GB memory.

5.2 Effectiveness (RQ1)

To answer RQ1, we assessed the effectiveness of LAFF to suggest appropriate values for each of the target fields in the dataset. We compared LAFF with MFM (most frequent model), ARM (association rule mining) [53], NaïveDT (naïve application of decision trees), and FLS (first letter search), which are able to provide suggestions under different form filling orders:

- (1) MFM is a widely-used form filling algorithm, which suggests possible values of a target field based on their frequency in historical input instances.
- (2) ARM is a state-of-the-art algorithm for form filling. ARM uses historical input instances to mine association rules with a minimal level of support and confidence; it matches the filled fields with mined association rules, and suggests the consequents of the matched rules to users.
- (3) NaïveDT is a naïve application of decision trees for form filling. We use decision trees because this type of model has been already used in the form filling literature [38]. Given a target field, this approach takes a subset of the remaining fields as filled fields (i.e., features); it then trains a decision tree for each feature-target combination. During form filling, based on the filled fields and the target field, NaïveDT selects the decision tree trained on the same feature-target combination in order to predict the target field. We considered this naïve application of decision trees because previous work [38] has shown that the effectiveness of a single decision tree trained for each target field is poor (see also section 7); our preliminary evaluation has also confirmed this.
- (4) FLS simulates form filling in categorical fields through a “typing” function. This function allows users to type the first letter of the candidate value they intend to fill (i.e., the first letter of the ground truth in this study). FLS filters the list of candidate values based on this letter and presents the refined candidate values as suggestions.

We did not consider other approaches for automated form filling, since they rely on additional information beyond the input values provided in the past for the same form. For example, they reuse the values filled in other software systems [9], extract information from text files (e.g., a CV file to fill job search sites) [69], or refactor forms for effective form filling [17]. An empirical comparison with these techniques is not feasible, since such additional knowledge is not always available during form filling; moreover, LAFF does not assume the existence of such knowledge. We discuss the differences between LAFF and these related approaches in section 7.

Choosing effectiveness metrics. We reviewed the main metrics used for evaluating the effectiveness of building a suggestion list. More specifically, we investigated the metrics used in the recommender systems area because, similar to automated form filling, many software applications use recommender systems to support software stakeholders in their decision-making while interacting with large information spaces [63] (e.g., locating faulty code snippets in software projects [81]). Table 2 shows, for each metric we reviewed, its dimension, description, and rationale.

⁵Due to the data protection policy of our partner, we were obliged to run the experiments on the PROP dataset using an on-premise, dedicated server that, however, could not be used to store external data (like the NCBI dataset).

Table 2. Main metrics used in the recommender system area for evaluating the effectiveness of building a list of suggestions

Dimension	Description	Intuition
Diversity	Diversity metrics generally measure the average dissimilarity between all pairs of items in the suggestion list [51].	The suggested items can cover a broad area of the information space to increase the chance of satisfying the user's information need. For example, with a movie recommender system, users may hope to get relevant items from different genres (e.g., "comedy", "romance")
Novelty	The novelty of an item is typically estimated by the inverse of its popularity (e.g., measured by the number of ratings it has received); novelty metrics measure the ratio of the suggested relevant items that have low popularity [46].	A tool has the ability to suggest relevant items that are unknown to users. For example, a movie recommender system should have the ability to recommend "new" movies that users did not watch or know before.
Accuracy	<i>Precision</i> measures the fraction of suggested relevant items among all the suggested items [37].	The suggestion list only contains relevant items.
	<i>Recall</i> measures the fraction of suggested relevant items among all the relevant items [37].	The suggestion list contains all the relevant items regarding a target.
	<i>MAP (Mean Average Precision)</i> measures the mean of the average precision at the rank of each relevant item [37].	All the relevant items can be ranked at the top of a suggestion list.
Coverage	<i>MRR (Mean Reciprocal Rank)</i> measures the mean of the reciprocal rank at which the first relevant item was suggested in a list [53].	The relevant item can be ranked at the top of a suggestion list.
	<i>Catalog coverage</i> measures the ratio of the items suggested to users over the total number of candidate items of a given target [30].	A tool can avoid suggesting a long list of items to users (e.g., only the top 5% suggested items are presented to the users)
Combined	<i>Prediction coverage</i> measures the ratio of suggestions provided by a tool over the total number of targets requiring suggestions [30].	A tool can avoid making "useless" suggestions to users. Low-confidence suggestions with many unrelated items should be filtered out to fit the user's interests.
	F_1 -score is the harmonic mean of the precision and recall [10].	A tool can suggest (recall) and only suggest (precision) all the relevant items to users.
	<i>Quality</i> of suggested items is the product between their similarity to the user's query and the diversity of the items [10].	The items suggested in the suggestion list are relevant to the user's query (similarity) and at the same time different from each other (diversity).

Metrics for evaluating recommender systems span over four dimensions, including diversity [51], novelty [16, 46], accuracy [37], and coverage [30, 37]. As shown in Table 2, diversity and novelty focus on the dissimilarity among the suggested items: the former by looking at pairwise dissimilarity, and the latter by determining the difference between the currently and previously suggested items. Both metrics can be applied in contexts where more than one relevant item can be suggested. As for assessing accuracy, precision and recall are the most common metrics [37]; they measure the ratio of correctly suggested items. However, precision and recall ignore the exact ranking of items as only the correct or incorrect classification is measured [67]. Other common accuracy metrics include MRR and MAP [37, 48, 67], which are designed to evaluate a list of suggested items. MRR calculates the rank of the first relevant item, and MAP measures the average precision of relevant items at different positions. Regarding coverage, two definitions have been proposed in the literature [30]: catalog coverage and prediction coverage. Catalog coverage measures the length of a list of suggested items relative to its maximum length; prediction coverage calculates the ratio of targets for which an algorithm provides suggestions over the total number of targets requiring suggestions. Finally, in the literature, some metrics are also proposed to combine different metrics from the same dimension to evaluate the trade-off between them. For example, metrics like F_1 -score and the quality metric [10] have been proposed to balance the weight of precision and recall (accuracy dimension) or similarity and diversity (diversity dimension), respectively.

According to a previous study [63], an effective recommender system in software engineering (e.g., a form filling system) is expected to avoid “helpless” suggestions that may be ignored by users, but provide users a large number of “helpful” suggestions. Considering the dimensions in Table 2, metrics for diversity and novelty are not applicable for form filling, as most categorical fields only contain a single correct value for a user to select. The accuracy metrics can be used to assess the “helpfulness” of suggestions. Since LAFF suggests a list of values for users to select the correct one, MRR is the most appropriate metric in our context, since it evaluates if the correct value is ranked at the top of a suggestion list. Regarding coverage, we select the prediction coverage to evaluate the extent to which the endorser module of LAFF can avoid “helpless” suggestions; it calculates the frequency of suggestions made by LAFF when required to make one. Since MRR and prediction coverage belong to different dimensions, we separately evaluate the two metrics instead of combining their results with a single score.

In the following, we provide the definition of MRR and prediction coverage.

MRR (Mean Reciprocal Rank) is defined as:

$$MRR = \frac{1}{|S|} \sum_{i=1}^S \frac{1}{k_i}, \quad (1)$$

where $|S|$ is the number of target fields that the algorithm provides suggestions, and k_i is the position of the first correct value in the i -th suggestion.

Prediction coverage rate is defined as:

$$PCR = \frac{|S|}{|S_{all}|}, \quad (2)$$

where $|S|$ is again the number of target fields that the algorithm provides suggestions, and $|S_{all}|$ is the total number of target fields in all the testing input instances [30, 37].

Methodology. To assess the effectiveness of the various form filling algorithms, we computed *MRR* and the prediction coverage rate *PCR*.

We remind the reader that LAFF uses the filled fields (i.e., features) of each test instance to predict the value of a target field. In our datasets, all the target categorical fields only have a single correct value (e.g., in Figure 6, *Leasing Serv.* is the ground truth for the field f_5 of the input instance #7).

For each test instance, we checked the position of the correctly suggested value (i.e., the value that corresponds to the ground truth) in each suggestion list and computed the reciprocal rank of the correct value. If no correct value was found, we set the reciprocal rank to zero. For example, in Figure 6, LAFF suggested three values; the reciprocal rank of the suggestion for ST_2 is 1 since the user can find the correct value to fill in first position (i.e., $k_i = 1$). The *MRR* value was computed as the mean of the reciprocal ranks for all the suggestion lists. Given a test set, we calculated the average *MRR* value for different targets. Concerning *PCR*, we counted the number of target fields for which an algorithm provided suggestions (i.e., a list of suggested values) and the number of target fields for which no suggestion was provided. *PCR* was computed as the percentage of target fields receiving suggestions over the total number of target fields.

As there is no publicly available implementation of the baselines (ARM, MFM, NaïveDT, and FLS) for form filling, we implemented them from scratch.

In the case of ARM, we set the minimum acceptable support and confidence to 5 and 0.3, respectively, consistent with previous work [53].

We implemented NaïveDT using the open-source library *scikit-learn* [60]. For the NCBI dataset, after the preprocessing step (see section 4.1) we obtained six fields; all of them are categorical. According to the discussion in section 4.2, we have $n = 6$ and $t = 6$, resulting in $6 * (2^{6-1} - 1) =$

Table 3. *MRR* and *PCR* of Form Filling Algorithms (t/o: timeout; N/A: not applicable)

Alg.	Sequential		Random		Train (s)	Predict (ms)		
	MRR	PCR	MRR	PCR		avg	min-max	
NCBI	MFM	0.42	1.00	0.42	1.00	0.03	0	0
	ARM	0.47	1.00	0.53	0.99	1.26	120	8-409
	NaïveDT	0.52	1.00	0.53	1.00	23	1	1-1
	FLS	0.55	1.00	0.54	1.00	N/A	0	0
	LAFF	0.74	0.70	0.73	0.70	546	15	5-44
PROP	MFM	0.64	1.00	0.64	1.00	0.15	0	0
	ARM	0.67	0.95	0.65	0.96	15	878	6-3074
	NaïveDT	t/o	t/o	t/o	t/o	t/o	t/o	t/o
	FLS	0.53	1.00	0.58	1.00	N/A	0	0
	LAFF	0.78	0.86	0.81	0.87	3652	138	9-317

186 feature-target combinations (i.e., decision trees to train). For the PROP dataset, after the preprocessing step, we obtained 15 fields, among which 14 are categorical (i.e., $n = 15$ and $t = 14$). This leads to $14 * (2^{15-1} - 1) = 229\,362$ decision trees to train.

Regarding FLS, we ranked the candidate values for a given target field alphabetically and refined these values by the first letter the user intends to fill. We assume that the first letter of the value of the ground truth is what the user intends to fill. Based on this assumption, FLS suggests the candidate values that start with the same letter as the ground truth.

We set a timeout of 24 hours to train each algorithm. This timeout value reflects the realistic situation in which the algorithm gets daily updates of its models, empowered with the information derived from new input instances collected throughout the day.

Results. Table 3 shows the effectiveness of the various algorithms for the two form filling scenarios. We remark that NaïveDT timed out during the training phase on the PROP dataset⁶

In terms of coverage rate (columns *PCR*), MFM, ARM, NaïveDT, and FLS provide suggestions on almost all the target fields ($PCR \approx 1$), while LAFF achieves a *PCR* value ranging from 0.70 to 0.87 on the two datasets. The lower *PCR* value of LAFF is ascribable to its endorsing module, which discards low-confidence suggestions.

As to the accuracy of these suggestions, MFM, ARM, and NaïveDT achieve *MRR* values ranging from 0.42 to 0.53 on the NCBI dataset; MFM and ARM achieve *MRR* values ranging from 0.64 to 0.67 on the PROP dataset. LAFF substantially outperforms MFM, ARM, and NaïveDT, achieving a *MRR* value above 0.73 in both datasets. The improvement in *MRR* value obtained by LAFF over NaïveDT is +22 pp on the NCBI dataset; the one over ARM is +11 pp on the PROP dataset for the sequential filling scenario. For the random filling scenario, the improvement over NaïveDT and ARM is +20 pp on the NCBI dataset; the improvement over ARM is +16 pp on the PROP dataset. According to Table 3, LAFF also outperforms the interaction-based approach FLS by +19 pp to +25 pp on the two datasets for different filling scenarios. The reason is that, after refining the candidate values by typing letters, users could still have dozens of candidate values with the same initial letter to check. For example, when users type “L” to find the country “Luxembourg”, the form returns 9 results (e.g., “Laos”, “Latvia”, “Lebanon”) where “Luxembourg” is the last one in the refined list; in contrast, LAFF can rank “Luxembourg” as the first country for users to choose, leading to much higher *MRR* values. Compared with FLS, LAFF has two advantages. First, many users lack a detailed conceptual

⁶For completeness, based on our preliminary evaluation, we note that using a single decision tree trained for each target field (as done in previous work [38]) would have avoided the time-out on the PROP dataset but would have yielded worse effectiveness results than the lowest values reported in Table 3, confirming the trend already reported by Hermens and Shlimmer [38] (see also section 7).

model of the software system [40] defined by the requirements analysts and domain experts. They may not remember all the candidate values predefined in a categorical field, leading to a potential lengthy search process [40]. In this case, LAFF can directly provide the most-likely suggestions for users to choose based on the filled fields. Second, LAFF is compatible with FLS. Based on the highly accurate suggestions made by LAFF, users can continue refining the suggested list with FLS when needed to further accelerate their form filling process. For example, when users type “L” after LAFF’s suggestion, only country names starting with “L” can be retained (e.g., “Luxembourg”, “Laos”, “Latvia”, “Lebanon”).

We use the Mann-Whitney U test to assess the statistical significance of the difference between the MRR values of LAFF and the baselines, with a level of significance $\alpha = 0.05$. The results show that LAFF always achieves a statistically higher MRR value than the baselines for the two form filling scenarios on both the NCBI and PROP datasets (p -value < 0.01).

These results have to be interpreted with the usage scenarios of a recommender system. Previous studies show that, for a recommender system, inaccurate suggestions increase users’ decision time and the risk of making wrong decisions [58]. The MRR and PCR values achieved by LAFF show that the suggestions provided by LAFF allow users to find the correct value among the top-ranked suggested values.

Error analysis. We further analyzed the suggestions made by LAFF, to identify the cases in which it does not perform well. We recall that in our experiments, LAFF suggested the top 5% most likely values for each target field. On the NCBI dataset, LAFF captures the correct value in the top 5% suggested values for 79.0% of the suggestions when using sequential filling and for 79.9% of the suggestions when using random filling. On the PROP dataset, the correct value is in the top 5% suggested values for 89.5% (sequential filling) and 90.1% (random filling) of the suggestions. Overall, 79.0% to 90.1% of the suggestions made by LAFF allow users to find the correct value among the 5% top-ranked suggested values. For the remaining incorrect suggestions (around 21% on the NCBI dataset and 9.9% on the PROP dataset), in which the correct value is not in the top 5% suggested values, we identified the following main reasons.

First, LAFF tends to provide incorrect suggestions when the number of filled fields used for prediction is small. Specifically, on the NCBI dataset, 33.1% of the incorrect suggestions when using sequential filling and 53.1% of the incorrect suggestions when using random filling were made when there was only one filled field; on the PROP dataset, the ratio is 17.4% and 23.5%, respectively. With few filled fields, LAFF may not get enough knowledge (i.e., the information of dependent field values) for prediction. This affects both the variable inference step within BNs and the behavior of the endorser module. In this case, LAFF tends to use the most frequent value in the target field for prediction, since this value has a higher prior probability. One possible way to mitigate this issue, to be investigated as part of future work, could be to define form refactoring techniques that allow users to first fill fields that provide additional knowledge used to predict the values of other fields. We analyze the impact of the number of filled fields on the effectiveness of LAFF as part of RQ4 (§ 5.5).

Second, incorrect suggestions are caused by the number of training input instances. Due to optional fields, users may not provide values for all the fields. For example, for the field “ethnicity” in the NCBI dataset, only 15.6% of input instances contain a non-empty value. The sparseness of the filled values for a field leads to a small number of training input instances to learn the corresponding dependency. Continuing the example, the MRR values for the field “ethnicity” are 0.608 and 0.553 for the sequential filling and random filling scenarios, respectively, thus leading to many incorrect suggestions. To mitigate this problem, as part of future work, instead of using a single threshold for endorsing suggestions, we could modify the endorser used in LAFF to support field-specific

thresholds. We analyze the impact of the size of the training set on the effectiveness of LAFF as part of RQ5 (§ 5.6).

Third, the number of options (i.e., candidate values) for a field *may* affect the effectiveness of LAFF. To investigate this, we computed the correlation between the number of options and the MRR value for a field, considering the MRR values achieved in the random filling scenario⁷. The resulting Pearson correlation coefficient is -0.09 on the NCBI dataset (p -value=0.722) and -0.477 on the PROP dataset (p -value=0.001), thus showing no correlation for NCBI and a moderate but significant correlation for PROP [23]. The difference in results between datasets can be easily explained by the fact that the variance in number of options is very low for NCBI (228.8), while it is much larger for PROP 6713.25. These results therefore suggest that, when a field has more options, LAFF tends to provide more incorrect suggestions.

To conclude, *the answer to RQ1 is that LAFF can yield a large number (with a PCR value ranging from 0.70 to 0.87) of accurate suggestions, with a MRR value above 0.73, significantly outperforming state-of-the-art approaches.*

5.3 Performance (RQ2)

To answer RQ2, we measured the execution time required to perform the model building phase of LAFF (i.e., training time), as well as the time to predict a target field (i.e., prediction time). The training time indicates the feasibility of using LAFF in contexts where the training set (i.e., the set of historical input instances) is updated often as new input instances are recorded in the system. The prediction time indicates how fast LAFF can provide form filling suggestions during a data entry session.

Methodology. We used the same settings (i.e., form filling scenarios) as in RQ1. We computed the training time as the time to build all BN models over the historical input instances. The prediction time is the average time (over the various target fields) taken to provide a suggestion for one input instance using locally deployed models. We also compared LAFF with the MFM, ARM, and NaïveDT algorithms. Notice that FLS does not require any training and the prediction can be considered instantaneous.

Results. The results are shown in the last two columns in Table 3, in term of training time (column *Train*) and the prediction time (column *Predict*, with sub-columns indicating the average, minimum, and maximum values, when applicable) for the two datasets.

The training time of LAFF is much higher than the one of MFM, ARM, FLS (and NaïveDT on the NCBI dataset); LAFF takes 546 s and 3652 s to train models on the NCBI and PROP datasets, respectively. This can be easily explained since LAFF trains several models, as explained in section 4.2. Although NaïveDT took only 23 s to train all the decision trees on the NCBI dataset, it timed out on the PROP dataset. As for prediction time, that of LAFF is higher than that of MFM, NaïveDT, and FLS. MFM and FLS directly suggest the frequency-based (for MFM) or matching-based (for FLS) value list to users. For NaïveDT, the prediction time includes the time to select the appropriate model (based on the feature-target combination) among the trained models. However, the prediction time of LAFF is, on average, faster than the one of ARM by 105 ms (15 ms vs 120 ms) over the NCBI dataset and by 740 ms (138 ms vs 878 ms), over the PROP dataset. The Mann-Whitney U test also confirms that the differences in prediction time between LAFF and the baselines are statistically significant (p -value < 0.01 for the two datasets).

⁷We did not consider the sequential filling scenario, since it could introduce some bias in our analysis: The number of filled fields to predict for each target field is different, as it depends on the tabindex order of the field.

Table 4. Effectiveness of LAFF with Different Modules

ID	Module		NCBI				PROP			
			Sequential		Random		Sequential		Random	
	L	E	MRR	PCR	MRR	PCR	MRR	PCR	MRR	PCR
LAFF-LE	✗	✗	0.54	1.00	0.56	1.00	0.74	1.00	0.77	1.00
LAFF-E	✓	✗	0.60	1.00	0.60	1.00	0.76	1.00	0.79	1.00
LAFF-L	✗	✓	0.55	0.83	0.61	0.77	0.78	0.86	0.79	0.87
LAFF	✓	✓	0.74	0.70	0.73	0.70	0.78	0.86	0.81	0.87

These results have to be interpreted taking into account the usage scenarios of our approach. The training time has to be considered when training the models with new data, i.e., the new input instances recorded in the system since the last execution of the model building phase. This task is performed *offline* and *periodically* (e.g., once a day), so a training time of the order of one hour is acceptable from a practical standpoint.

Given the interactive nature of data-entry applications, having a short prediction time is much more important for an automated form filling approach. According to human-computer interaction principles [36], users feel a system reacts instantaneously when its response time is within 100 ms and feel they are seamlessly interacting with the software system when the response time is within 1 s. In our context, the prediction time of LAFF depends on the computational power at our disposal and the complexity of the trained BN models (i.e., the number of nodes and the size of probability tables). The experiments using the proprietary dataset (extracted from a production-grade system) show that LAFF is fast enough (requiring at most 317 ms) to provide real-time suggestions during data entry sessions.

The answer to RQ2 is that the performance of LAFF, with a training time of about one hour (or less) and a prediction time of at most 317 ms, is suitable for practical application in data-entry scenarios.

5.4 Impact of Local Modeling and Endorser (RQ3)

LAFF has two important modules: the local modeling module, which builds local models based on local field dependencies of partitions of historical input instances (section 4.2); the endorsing module, which uses heuristic rules to remove possibly inaccurate suggestions (section 4.3). To answer RQ3, we assessed the impact of these two modules on the effectiveness of LAFF.

Methodology. As shown in Table 4, in the two sub-columns of column *Module*, we considered four variants of LAFF, to reflect possible configurations with the two modules. In the table, *L* refers to the local modeling module and *E* refers to the endorsing module; symbols ‘✓’ and ‘✗’ indicate whether the variant of LAFF includes or not a certain module, respectively. When the local modeling module is disabled (denoted by LAFF-L), it means LAFF only uses the global model for prediction; when the endorsing module is disabled (denoted by LAFF-E), it means we do not scrutinize (and possibly discard) the suggestions provided by LAFF. If both modules are disabled (denoted by LAFF-LE), LAFF becomes a plain BN model trained on the entire set of historical input instances. We ran the vanilla version of LAFF (i.e., the one presented in section 4) and the additional variants using the same settings as in RQ1, and measured effectiveness in terms of *MRR* and *PCR*.

Results. As shown in Table 4, each module impacts the effectiveness of LAFF. The local modeling module improves the ability of BNs in ranking the correct values ahead of the incorrect ones, leading to a higher *MRR* value (while the *PCR* value remains equal to 1). The endorsing module mainly reduces the quantity of inaccurate suggestions made by different BN models (and therefore reduces the *PCR* value).

When we compare LAFF (with both modules enabled) with a plain BN (i.e., LAFF-LE) on the NCBI dataset, LAFF improves the *MRR* value by +20 pp (0.74 vs 0.54) for the sequential filling scenario and by +17 pp (0.73 vs 0.56) for the random filling scenario; on the PROP dataset the improvement is smaller (+4 pp for both scenarios). Hence, the integration of the local modeling and endorsing modules positively affect the effectiveness of LAFF. When we apply the endorsing module on a plain BN, we get an *MRR* improvement ranging from +1 pp to +5 pp on the two datasets (see LAFF-LE vs LAFF-L). Even though the local modeling module alone does not affect the number of suggestions, the integration of the local modeling in LAFF-L leads to a further reduction of the number of inaccurate suggestions (with *PCR* dropping by 0 pp to 13 pp); it improves the *MRR* value by 0 pp to +19 pp on the two datasets (see LAFF-L vs LAFF).

These results can be explained as follows. As mentioned in section 4.3, the endorsing module endorses suggestions based on two heuristics: *checkDep* (which checks if the filled fields are parents of the target field) and *sumProb* (which checks whether the sum of probabilities for the top- n_r suggested values is higher than a threshold). When the local modeling module is enabled, the local models can learn fine-grained dependencies for a field and exclude some useless dependencies that are present in a plain BN. As a result of the absence of the dependencies between filled fields and the target field, the *checkDep* heuristic may endorse more suggestions (i.e., the *PCR* value decreases from LAFF-L to LAFF), in order to retain high-confidence suggestions (i.e., the *MRR* value increases). On the PROP dataset, the improvement is not obvious because of the high quality of its input instances. As a proprietary dataset from the banking domain, the data quality division usually double-checks the data entry to minimize the effect of data errors and data conflicts on financial software systems. Compared to a public dataset, BNs trained on the PROP dataset can find more meaningful field dependencies and also achieve high *sumProb* values for most suggestions; hence the differences in both *MRR* and *PCR* values are relatively small when different modules are enabled.

The answer to RQ3 is that the local modeling module and the endorsing module improve the effectiveness of LAFF.

5.5 Impact of the Number of Filled Fields (RQ4)

LAFF takes as input a set of already-filled fields with their values to suggest the value of the target field. To answer RQ4, we assessed the impact of the number of filled fields on the effectiveness of LAFF as well as of the ARM and MFM baselines. We did not compare to the NaïveDT and FLS baselines. NaïveDT is impractical to be used in a production-grade system due to the timeout issue during the training phase. FLS does not use the information contained in the already-filled fields.

Methodology. We generated new test sets by varying the number of filled fields on the testing input instances obtained as described in section 5.1. To generate a test set with i filled fields for a target t , for each testing input instance, we set the field t as the target and randomly selected i non-empty fields as filled fields. The unselected fields were considered as unfilled and their values replaced with a dummy value representing empty fields. Given a data entry form with t targets for automated form filling, we can generate t new test sets with i filled fields, each of which has a different target. We ran LAFF on these new test sets and computed the *MRR* and *PCR* values for predicting different targets. The results indicate the effectiveness of LAFF when i fields are filled.

Following this strategy, we assessed the effectiveness of LAFF and the baselines on the NCBI dataset with one, two, and three filled fields. We discarded the configuration with four filled fields, since we could only generate 505 new testing input instances due to the optional fields; this number is significantly smaller than the number of testing input instances we obtained for the configurations

with one/two/three filled fields (which have more than 17 000 new testing input instances) and might have introduced bias.

For the PROP dataset, we generated 16 test sets representing the configurations with one to 16 filled fields; each of them has more than 300 000 testing input instances. However, running the experiment for all testing input instances would be infeasible on the dedicated server provided by our industrial partner, which caps the duration of any job to 168 hours. The issue of the execution time was mainly introduced when evaluating ARM. As mentioned in section 5.3, ARM may take more than 3 s to provide a suggestion. This means ARM could take, in the worst case, $\approx 300\,000 \times 16 \times 3\text{ s} \approx 166$ days to execute on the testing input instances for all the 16 generated test sets in the PROP dataset. Hence, to assess the impact of the number of filled fields on all the algorithms when using the PROP dataset, we randomly sampled 12 600 testing input instances from each of the 16 newly generated test sets. We chose this number because, in the worst case, the experiments for all the algorithms could be finished within the 168 hours limit (e.g., $12\,600 \times 16 \times 3\text{ s} = 168$ hours for ARM).

Results. Figure 7 shows the results of running the form filling algorithms on the NCBI and PROP datasets with different numbers of filled fields. The x-axis of the figures represents the number of filled fields and the y-axis shows the *MRR* or *PCR* value of each form filling algorithm.

In terms of accuracy (in Figure 7(a) and Figure 7(c)), the *MRR* values of LAFF and ARM increase when more fields are filled. For example, the average *MRR* value of LAFF increases from 0.599 ± 0.227 to 0.730 ± 0.350 on the NCBI dataset and from 0.774 ± 0.160 to 0.804 ± 0.190 on the PROP dataset as the number of filled fields increases from one to three and from one to 16, respectively. In contrast, MFM shows a different trend. In theory, MFM is not affected by the number of filled fields, as it always provides the same suggestion for a target field, based on the most frequent values filled in the target field in the past. However, our experiments show a variation in *MRR* as the number of fields increases. More specifically, since the testing input instances generated for each number of filled fields are different, the *MRR* value of MFM decreases when the correct values for a target in the generated testing input instances are not the most frequent historical values. Overall, according to the boxplots, when varying the number of filled fields, the *MRR* values of LAFF for predicting different targets remain higher than those of MFM and ARM. Given a partially filled form with only one filled field, LAFF outperforms MFM and ARM by +12 pp and +11 pp, respectively, on the NCBI dataset; it also outperforms both MFM and ARM by +14 pp on the PROP dataset.

As for coverage, the *PCR* value of the baselines is 1 for the majority of the numbers of filled fields (as shown in Figure 7(b) and Figure 7(d)). For LAFF, the *PCR* value increases as more fields are filled. However, we find that the *PCR* values of LAFF for different target fields vary significantly, especially with a small number of filled fields. For example, with one filled field, the *PCR* values of LAFF are 0.482 ± 0.22 and 0.630 ± 0.41 respectively on the NCBI and PROP datasets. This is because, when the number of filled fields is small, the filled fields may not always provide enough knowledge (i.e., the information of dependent field values) for LAFF to predict all the target fields. For some targets, the endorsing module of LAFF may filter out many suggestions. When the number of filled fields increases, LAFF gets more information from the input instances to provide suggestions that can be endorsed. For example, the *PCR* value of LAFF increases to 0.975 ± 0.04 on the PROP dataset with 16 filled fields. Overall, The *PCR* values show that LAFF could correctly endorse the suggestions when the form is partially filled, which helps LAFF achieve higher *MRR* values than those of the baselines.

The answer to RQ4 is that the effectiveness (in terms of MRR and PCR) of LAFF increases as more fields are filled. Further, LAFF can better handle partially filled forms than state-of-the-art algorithms.

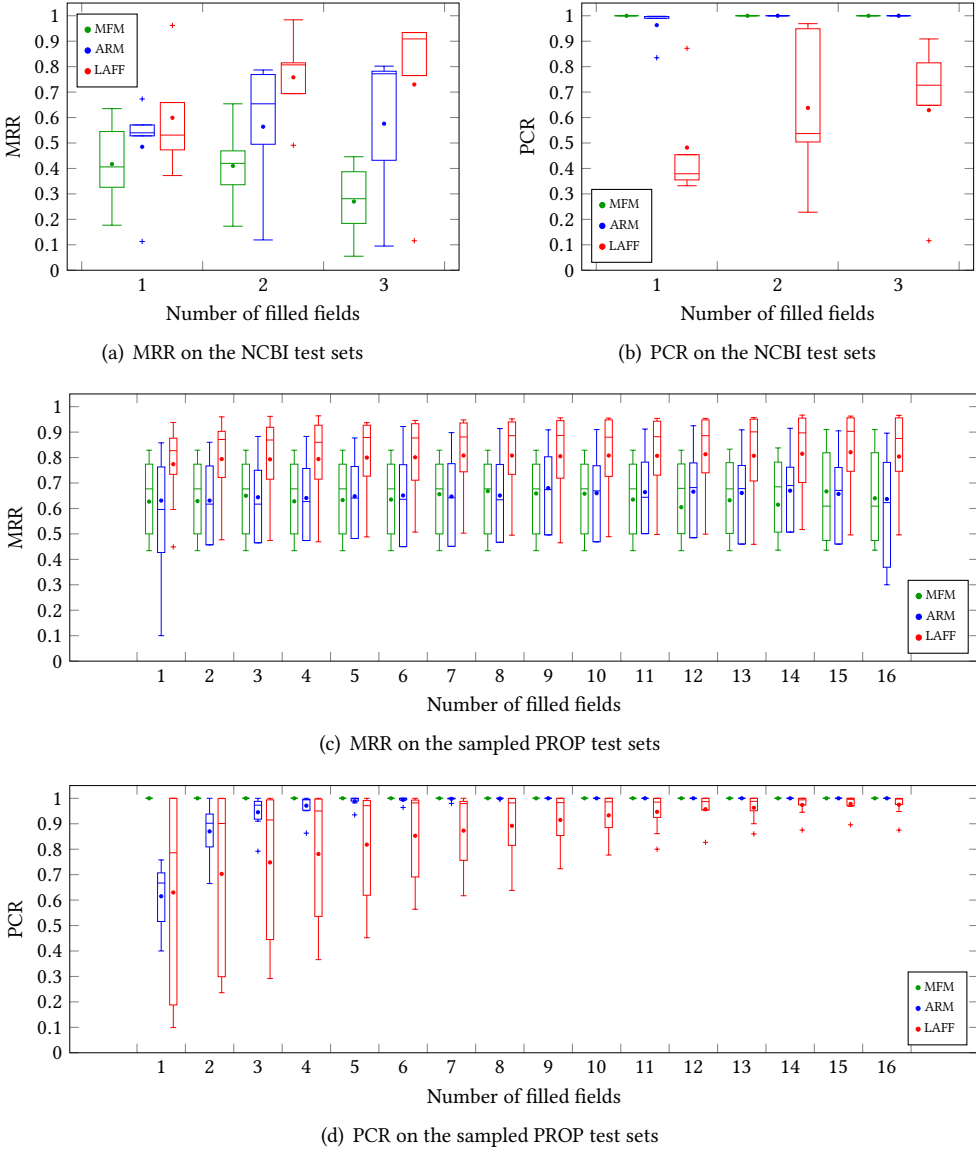


Fig. 7. Effectiveness of LAFF with different number of filled fields on the NCBI and PROP datasets

5.6 Impact of the Size of the Training Set (RQ5)

LAFF is a learning-based approach that requires a training set (i.e., historical input instances) to train machine learning models. To answer RQ5, we assessed the impact of the size of the training set on the effectiveness of LAFF.

Methodology. We evaluated LAFF by varying the size of the training set from 10% to 100% of the historical input instances included in the training set, with a step of 10%. For the NCBI dataset, the size of the sampled training set ranged from 5928 to 59284 historical input instances (where 59284

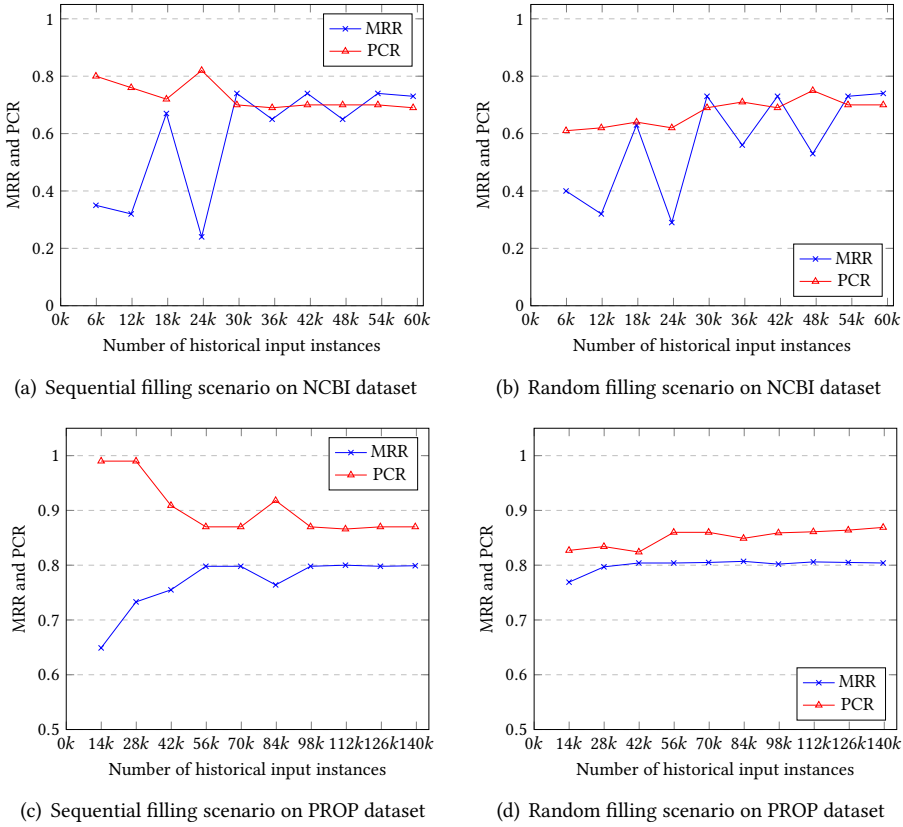


Fig. 8. Effectiveness of LAFF with different training set sizes on the NCBI and PROP datasets

is 80% of the 74 105 input instances in the dataset). For the PROP dataset, the size of the sampled training set ranged from 13 955 to 139 557 historical input instances (where 139 557 is 80% of the 174 446 input instances in the dataset).

For a given percentage value p , we randomly sampled $p\%$ of the historical input instances in the training set to form a smaller training set. We trained LAFF on the sampled training set and used the trained model to conduct automated form filling on the testing input instances obtained as described in section 5.1.

As for RQ1, we measured the effectiveness of LAFF in terms of *MRR* and *PCR*.

Results. Figure 8 shows the results of LAFF on the NCBI and PROP datasets with different training set sizes. The x-axis of the figure represents the number of historical input instances; the y-axis shows the *MRR* and *PCR* values of LAFF under different filling scenarios.

As shown in the figure, the value of *MRR* increases and gradually becomes stable on the two datasets when the size of the training set increases. For the NCBI dataset, the value of *MRR* is low and fluctuates significantly with a training set of less than 30 000 historical input instances; the majority of *MRR* values are lower than 0.40 for the two form filling scenarios. When we have more historical input instances (between 30 000 and 60 000), the value of *MRR* increases and becomes more stable, ranging from 0.650 to 0.740 for the sequential filling scenario and from 0.530 to 0.740 for

the random filling scenario. The value of MRR has a similar trend on the PROP dataset; it gradually increases and then becomes stable (between 0.764 and 0.807) when the number of historical input instances is higher than 56 000.

In terms of PCR , the endorser module of LAFF works poorly with a small training set. LAFF either keeps the majority of suggestions (for the sequential filling scenario) or wrongly removes many suggestions (for the random filling scenario), leading to low MRR values. As the size of the training set increases to more than 30 000 (for the NCBI dataset) and 56 000 (for the PROP dataset), LAFF is able to remove more inaccurate suggestions, achieving a PCR value between 0.690 and 0.750 and between 0.849 and 0.918, respectively on the two datasets.

When comparing the results of LAFF across the two datasets, we observe a more significant fluctuation of MRR values on the NCBI dataset than that on the PROP dataset as the size of the training set increases. This is caused by the data quality of the NCBI dataset. In contrast to the proprietary dataset PROP, there is no field constraint or additional check on the values filled in the NCBI data entry form. This means that, when more historical input instances are added to the training set, one may also introduce many conflicting or erroneous field values, which increase the uncertainty of LAFF when predicting on the NCBI dataset.

The answer to RQ5 is that the size of the training set affects the effectiveness of LAFF. MRR values increase on both datasets when the size of the training set increases; more suggestions are also correctly endorsed. With more than 56 000 historical input instances, LAFF achieves accurate suggestions on both datasets

5.7 Threats to Validity

The size of the pool of historical input instances can affect the effectiveness of LAFF, a common issue among learning algorithms. Nevertheless, we do not expect this to be a strong limitation since it targets data entry functionalities in enterprise software, in which one can expect thousands of input instances per day, as it is the case for the system used by our industrial partner.

Another threat to the validity is the choice of the value of the endorser threshold θ . With a higher threshold, LAFF will filter out more suggestions (resulting in a lower PCR value), only keeping the ones with a high predicted probability (resulting in higher MRR value). Hence, this threshold reflects the degree of uncertainty users are willing to accept regarding the suggestions provided by LAFF. To mitigate this threat, we selected a threshold value based on the feedback received by data entry operators and data quality engineers of our partner.

The choice of the deployment method for LAFF can impact its performance in terms of prediction time. In our experiments, we deployed LAFF locally; using a different deployment (e.g., cloud-based) could lead to different results, since the prediction time would be affected by many other factors, such as the DNS lookup time, the connection time, and the data transmission time [15]. Since the prediction time of LAFF is less than 317 ms, an application using a non-local deployment would have enough leeway to optimize these factors and provide seamless interactions for users, complying with human-computer interaction principles [36] (i.e., with a response time less than 1 s). As part of future work, we plan to assess the performance of LAFF under different deployment configurations.

To increase the external validity of our results, LAFF should be further studied on other datasets, possibly from other domains. To partially mitigate this threat, we use two different types of datasets to evaluate LAFF, including both a public and a proprietary dataset, and the corresponding data entry forms. Meanwhile, we simulated two form filling scenarios (sequential and random filling) that are plausible during a real data-entry session. As part of future work, we plan to conduct a user study using different datasets and data entry forms, to analyze the effect of LAFF on reducing form filling time and input errors. Another external threat is our implementation of the four algorithms (MFM, ARM, NaïveDT, and FLS) to which we compared, which may be different from the original

definitions; to mitigate this threat, two of the authors cross-reviewed the implementations, taking into account the relevant literature (when available).

5.8 Data Availability

The implementation of LAFF, the NCBI dataset, and the scripts used for the evaluation are available at <https://figshare.com/s/b191fcb5ee7f3ad634bf>; LAFF is distributed under the MIT license. The PROP dataset cannot be distributed due to an NDA.

6 DISCUSSION

6.1 Usefulness

The fundamental question we seek to answer is whether LAFF can help users fill forms. To answer this question, we evaluated LAFF with two real-world form filling datasets, one from the biomedical domain and another from the banking domain. The results show that LAFF outperforms state-of-the-art form filling algorithms in providing a larger number (with a *PCR* value over 0.70) of accurate suggestions (with a *MRR* value over 0.73). The *MRR* value reflects the ability of LAFF in avoiding inaccurate suggestions. For example, considering a list with three suggested values, if the correct value is in the top-1, top-2, and top-3 of the list, the corresponding *MRR* value is 1, 0.5, and 0.33, respectively; when all the values are incorrect, the *MRR* value is 0. In the context of our experiments, an *MRR* value of at least 0.73 indicates that, for more than 73% of suggestions, LAFF can help users find the correct value from the top-ranked ones. The *PCR* value indicates the number of suggestions made by LAFF. A *PCR* value over 0.70 indicates that LAFF can confidently make suggestions for more than 70% of target fields, where the correct values are usually ranked before the incorrect ones. According to a previous study [17], as users are presented with more candidate values for selection, they tend to make more mistakes, which makes form filling a frustrating activity. Hence, we speculate LAFF reduces the mental load of users in filling forms by helping them go through fewer candidate values (top-k%) before finding the correct one; further user studies are required to corroborate this hypothesis.

We can interpret the above results from a point of view of usefulness as follows. On the one hand, as shown in RQ1, with an *MRR* value over 0.73, 79.0% to 90.1% of the suggestions made by LAFF allow users to find the correct value among the 5% top-ranked suggested values on the two datasets. This means that when LAFF makes suggestions, in 79.0% to 90.1% of the cases, it leads to at least 95% effort saving when browsing possible values, since users need only to check the top 5% most likely items recommended by LAFF. For example, on our datasets, users need to check between one (for field “activity” having 13 candidate values) and ten (for field “country” having 206 candidate values) values before finding the correct one. On the other hand, according to our analysis in RQ1, a frequency-based method like MFM achieves *MRR* values ranging from 0.42 to 0.64 on the NCBI and PROP datasets. Hence, when comparing LAFF with the widely-used data entry solution MFM, an *MRR* improvement of +14 pp to +32 pp suggests that LAFF can better help users select candidate values in data entry forms.

6.2 Practical Implications

This subsection discusses the practical implications of LAFF for its different stakeholders: software developers, end users, system administrators, and researchers.

6.2.1 Software Developers. Automated form filling is a common requirement for software systems. Some popular languages and APIs (e.g., HTML [74] and Android APIs [7]) also provide pre-fill or auto-completion frameworks, for which customized form filling strategies can be implemented. In its current version, LAFF is a stand-alone tool that developers can integrate into their data entry form

implementations as an effective and efficient strategy for filling categorical fields (e.g., Listboxes and Dropdown Lists). Since LAFF shows a higher accuracy than state-of-the-art approaches, it can be used in data-reliant enterprise systems across different domains, especially when there are constraints on sharing or accessing data of other software systems.

As for adopting LAFF in production, in terms of execution time, a training time of about one hour allows LAFF to compute daily updates for its models, empowered with the information derived from new input instances. Moreover, considering the interactive nature of data-entry applications, we use a model selection strategy to select the suitable local or global model for prediction from multiple models. Compared with many learning strategies that consider the predictions of multiple models jointly (e.g., ensemble learning), our strategy could significantly reduce prediction time in specific deployment configurations. For example, the worst-case prediction time of ensemble learning is, when using a single-thread configuration, the sum of the prediction time of all local models. This could be reduced to the maximum of the prediction time of all local models by using multiple threads in parallel. This strategy is viable when the trained models are deployed on a powerful server, and the prediction is performed server-side. However, some enterprises (as it is the case for our industrial partner) might prefer to integrate a form filling approach like LAFF *locally*, on the machines used by data entry operators, to avoid network latency during data entry sessions. In such a deployment configuration, using multiple threads could significantly slow down the data entry session, especially on machines with a limited computing power. For these reasons, we designed LAFF with a lightweight model selection strategy, instead of using alternatives like ensemble learning.

Integrating LAFF into a data entry form requires providing a mapping between the field names and the column names in the dataset. This mapping needs to be identified only once and therefore does not have much impact on the practical application of LAFF. We remark that, in modern applications built using the MVC pattern for the UI and an Object-Relational Mapping (ORM) framework for data persistence [18], the mapping can be retrieved from the ORM framework configuration. Other sources for the mapping can be software design documentation, such as the database schema and the description of the UI widgets in the data entry forms.

6.2.2 End Users. Form filling is time-consuming and error-prone for end users, which can cause more than half of data errors [61]. These errors seriously affect data-reliant software systems and even cause loss of human life [5, 11, 49, 61]. In this study, we have proposed LAFF to improve the accuracy and efficiency of the data entry process executed by end users, when filling categorical fields. First, LAFF uses an endorser module to alleviate the cognitive load on users caused by wrong suggestions; this module significantly improves the accuracy of suggestions. Second, LAFF helps users focus on the most-likely candidate values in a list of values. It decreases the number of candidate values users need to browse. *We follow the practice of designing an effective recommender system* (e.g., a form filling system) in software engineering, which is to avoid “helpless” suggestions to be ignored by users, but provide them with a large number of “helpful” suggestions [63].

Furthermore, in the experiments, LAFF can provide suggestions within at most 317 ms for each target field, which enables end users to conduct seamless interaction with the data entry form.

6.2.3 System Administrators. The deployment of LAFF in production requires system administrators to configure its parameters. This can be achieved with the help of domain experts, based on their domain knowledge. A group of configurable parameters is set in the pre-processing step. We implement this step based on best practices for predictive data mining [3]. A threshold that can affect auto-filling effectiveness is θ , for the endorser module. This threshold reflects how much uncertainty domain experts are willing to accept regarding the suggestions provided by LAFF. Such

configuration allows domain experts to use LAFF according to their requirements and application scenarios.

Overall, the configuration parameters represent the only domain-specific aspect of LAFF. Everything else about it is domain-agnostic. Deploying LAFF for a new form only requires to collect the input instances of the form for a certain period of time and use these instances to train the form filling models.

6.2.4 Researchers. In this paper, we use a local modeling module to effectively learn the fine-grained dependencies on historical input instances. Since existing approaches in software engineering perform local modeling on numerical data instances (e.g., software metrics [55]), we propose a novel solution to solve the problem of using local modeling on categorical data instances. We speculate that our proposed solution can inspire the adoption of local modeling for different data types in many software engineering tasks. In addition, we use an endorser module to decide if suggestions are accurate enough to be provided to end users. Such a module is important for algorithms where 100% accuracy cannot be achieved in practice. For example, during form filling, predicting all the other fields based on only one filled field may lead to inaccurate suggestions. In this case, it is more practical to automatically remove inaccurate suggestions using an endorser module. We believe the endorser-based architecture discussed in section 4.3 can be adopted by other recommender systems.

Furthermore, the error analysis for RQ1 described in section 5.2 suggests possible research directions, such as learning from users' corrections and using multi-objective optimization for form refactoring (with a new order of fields) for improving form filling suggestions.

6.3 Limitations

6.3.1 Type of Fields. In this work, we have focused on predicting the values of categorical fields, based on the historical information available from the same software system; LAFF does not work with other types of fields. Due to the unique characteristics of each type of field, distinct solutions have been proposed in the literature (e.g., text auto-completion for textual fields, data testing for numerical fields). However, these solutions cannot help users fill categorical fields, which is a critical task during form filling. We acknowledge that other types of fields may contain critical data that could cause major problems. However, as discussed in section 1, empirical studies show that selection errors lead to more than half (54.5%) of the data errors in a software system [61]. For example, the selection of wrong drugs [49] or the wrong modality of care [61] in medical record systems can even cause loss of human life. In addition, as shown in the NCBI and PROP datasets (in section 5.1), the majority of categorical fields we evaluated (with the number of candidate values ranging from 13 to 208) are related to certain domains or business processes. These fields are more difficult to fill than fields such as "sex" and "age", since users need to understand the meaning of candidate values. Errors in these fields can cause significant business problems. For example, the selection of a wrong "field of activity" when opening a bank account may cause business loss between the company and the bank. By knowing the actual "field of activity" of the company, the bank could have offered targeted products to its customer since the beginning of the business relation. Hence, there is a need for a semi-automated method that supports and guides users when filling categorical fields.

6.3.2 Cognitive Load on Users. The suggestions made by automated form filling tools like LAFF can increase the cognitive load on users, when the tool provides a long list of suggested values/options for users to check or when the suggestions do not include the correct value. To reduce the cognitive load in the first case, we configured LAFF to suggest the top 5% most likely candidate values for each target field, rather than reordering all the candidate values based on their probability.

Although LAFF does not directly reduce the number of candidate values in a field, it highlights the top 5% values, on which users can focus. As for the second case of cognitive load (suggestions not including the correct value), LAFF includes an endorser module to decide whether the suggestions are accurate enough to be returned to users. As shown in Table 4 and as part of the answer to RQ3, the endorser module significantly reduces the number of possibly inaccurate suggestions while increasing the *MRR* values. As a result, our experiments show that 79.0 % to 90.1 % of the LAFF suggestions allow users to find the correct value among the top 5% suggested values.

6.3.3 Fields with Semantic Overlap. Candidate values in categorical fields could have semantic overlaps. For example, let us consider the case of field “field of activity” with two possible values, “banking service” and “financial services”. From a semantic point of view, the former is a specific case of the latter. This form of semantic overlap affects LAFF as follows. During the model building phase, LAFF could be trained with inconsistent historical input instances (e.g., two historical instances that have same values in all fields but “field of activity”, with one instance having “banking service” and the other “financial services”). In this case, LAFF cannot build, with enough confidence, dependencies between the values in the field affected by a semantic overlap (e.g., “field of activity”) and the other fields. During the form filling suggestion phase, when LAFF identifies both values as candidate values to suggest (e.g., with similar probability), due to the endorser module, LAFF may not provide any suggestions or suggest both values to users (depending on the number of suggested values n_r and the threshold θ). In both cases, users need to decide which value to select by themselves. However, defining candidate values with semantic overlaps is not a good practice for form design, since it increases the mental load on users to decide which value is more appropriate (e.g., the more specialized “banking service” or the more generic “financial service”). We suggest to address this issue by refining the candidate values during the design phase of the system.

6.3.4 Cases of Limited Accuracy. As discussed in the “Error analysis” part of the answer to RQ1 (page 23), there are 9.9 % to 21 % of suggestions made by LAFF for which the correct value is not in the top 5% suggested values. We have identified two main reasons leading to low accuracy: LAFF tends to provide incorrect suggestions, when (1) the number of filled fields used for prediction is small, and (2) the size of the training input instances for a target field is small. We plan to address these limitations as part of future work, along the lines mentioned on page 23.

6.3.5 Learning from Corrections. Automated form filling is an interactive process between users and the automated tool. In the current version, LAFF does not offer the possibility, during a data entry session, for a user to correct LAFF’s suggestions by selecting a candidate value that is not presented in the top 5% suggested values. If a user makes many corrections to suggestions during a data entry session, it means that the learned probability table (BN) may not reflect the relations of values in the current input instance. The ability to learn from these users’ corrections could further improve the accuracy of LAFF. One intuitive solution is to assign a higher weight on such input instances when re-training LAFF for learning new relations. For example, input instances with many inaccurate suggestions can be oversampled to increase their proportion in the entire historical of input instances.

6.3.6 Cold Start. LAFF depends on the dataset associated with an input form; it always has to be trained for a new system. The size of the pool of historical input instances can affect the accuracy of LAFF, a common issue among learning algorithms. When there are no or only a few historical input instances (i.e., the cold start problem [52]), the accuracy of LAFF is limited. As presented in RQ5 (in section 5.6), *MRR* values of LAFF increase when the size of the training set increases; LAFF achieves accurate suggestions with 30 000 and 56 000 historical input instances on the NCBI

and PROP datasets, respectively. Nevertheless, we do not expect the size of the training set to be a strong limitation on the feasibility of applying LAFF. First, LAFF targets data entry functionalities in enterprise software, in which one can expect thousands of input instances per day. For example, nowadays the NCBI platform gets about 9600 input instances per month for a single species (i.e., “Homo sapiens”). It is also the case for the system used by our industrial partner. Second, LAFF does not require additional effort to label training data; it uses the actual values filled by users as the ground truth to train the model. With an adequate training set size, LAFF can directly provide accurate suggestions for users.

7 RELATED WORK

The approach proposed in this paper is mainly related to works on automated form filling based on the information from *the same software system*, which focuses on filling free-text fields and categorical fields. Regarding the former, the main proposals use language models (e.g., n-gram and sequence-to-sequence learning) to learn relationships between characters or words from historical textual inputs [66, 73]; these relationships are then used to provide word auto-completion based on the letters typed in a field [82]. As for dealing with categorical fields, most of the approaches suggest possible values from a list of candidate values. Martínez-Romero et al. [53] use association rule mining (ARM) to uncover the hidden associations of fields for real-time form filling; however, as shown in our experiments, ARM does not provide accurate suggestions when compared with LAFF. Hermens and Shlimmer [38] applied decision tree and hierarchical clustering for filling forms in an electronic leave report system: the study reports that these algorithms performed worse than the simple most-frequent method when the forms were filled out in a random order. Troiano et al. [71] trained a Bayesian network model on the historical inputs from a user for an online payment system. The trained model could auto-fill the payment form for the same user by reusing his historical payment records. The plain Bayesian network model is also trained by Ali and Meek [4] for auto-filling the bug submission form of a bug tracking system. Compared to existing approaches, LAFF performs automated form filling by mining field dependencies on input instances from different users. We proposed local modeling and a heuristic-based endorser to further improve the accuracy of form filling suggestions.

Automated form filling has also been investigated in the context of developing “smart” personal information management systems, to support information exchange *across software systems* [9, 19, 25, 79]. In this context, the main challenge is the semantic mapping of fields across software systems, e.g., how to map the “postal code” and “zip code” fields (from two different software systems) to the same concept. To address this challenge, Chusho et al. [19] manually construct rules to merge similar concepts of commonly used field names. Other works use string-based matching [35], WordNet [9], and Wikipedia [34] as additional resources to calculate the similarity among fields. Wang et al. [76] calculate field similarity based on the field names, form topics, and names of neighbor fields. They recently propose to use learning-to-rank algorithms to further improve form filling effectiveness [77]. However, to learn mapping rules, these algorithms require access to the users’ personal records from different software systems. In contrast, we only use input instances from the targeted software system; thus, our approach can be used when there are legal or security constraints on sharing records across systems [79]. Moreover, the above approaches perform well only on common fields like “age” and “address”, and cannot cope with fields that are domain-specific and used only in few software systems.

The filling order of fields influences the ability of form filling algorithms [38]. Several works refactor data entry forms to provide effective supports for form filling, for example using fields dependencies [17, 71] or user roles [2]. All these approaches identify the fields that a target field directly depends on and change their (field) order so that they can be filled before the target field,

Table 5. Overview of related work

Scenario	Reference	Task/Example	Difference with LAFF
Form filling under the same software system	Salama et al. [66] Van Den Bosch and Bogers [73] Zhang et al. [82]	[Task] Predict or auto-complete the next character or word for textual inputs. [Example] They suggest <i>Angeles</i> when users type the word <i>Los</i> .	They build language models on the values in <i>textual fields</i> instead of categorical fields; the user's initial input is required.
	Martínez-Romero et al. [53] Hermens and Shlimmer [38] Troiano et al. [71] Ali and Meek [4]	[Task] Suggest the correct value from a list of candidate values in <i>categorical fields</i> . [Example] They suggest <i>Leasing Service</i> from a list of options in the field "primary field of activity".	We compared LAFF with algorithms MFM [38], ARM [53], and NaïveDT [38] as part of RQ1; we compared to a plain BN (equivalent to the one proposed in [71] and [4]) as part of RQ3.
Form filling across software systems	Araujo et al. [9] Chusho et al. [19] Firmenich et al. [25] Winckler et al. [79] He et al. [35] Hartmann and Muhlhauser [34] Wang et al. [76] Wang et al. [77]	[Task] Prefill form fields by reusing the values filled in other web forms; they build the mapping or ontology of field names across web forms. [Example] They use the value filled in "postal code" in one web form to help the same user fill "zip code" in another web form.	These approaches cannot be applied to fields that are domain-specific and used only in few software systems (e.g., medical systems); enterprise information system may have constraints on visiting or sharing records across systems. LAFF is designed to work in these scenarios.
Form refactoring	Chen et al. [17] Troiano et al. [71] Akiki et al. [2]	[Task] Re-order fields to support effective form filling. [Example] They refactor the field "role" as the first field, if they find "role" is informative to predict values of other fields.	It is a preliminary step, complementary to automated form filling
Software artifacts information auto-completion	Umer et al. [72] Burgueño et al. [14]	[Task] Predict values of certain fields in software systems. [Example] They predict "bug priority" based on the description of bug reports.	It relies on textual data (e.g., bug reports and requirement documents) for suggestions. LAFF does not rely on such information.
Data crawling	Hernández et al. [39] Kantorski et al. [47]	[Task] Crawl data from databases by form filling. [Example] They generate value combinations for fields in a job search site to crawl all the job information hidden in the database.	They aim to generate valid input values for data crawling, instead of helping users find the correct candidate value.
Data mining	Diaz et al. [22] Kristjánsson et al. [50] Toda et al. [69]	[Task] Extract data values from data-rich text to fill forms. [Example] They use a resume text file to fill several fields of forms in job search sites.	They rely on the extraction of information in data-rich files (e.g., text files and spreadsheets). LAFF does not rely on such information.

to increase the accuracy of predicting the latter. Form refactoring can be regarded as a preliminary step to our proposed approach.

Form filling can be considered as a task to auto-complete data entry for software systems, which in the literature is a popular research topic in many domains; however, most of the existing approaches propose auto-completion for different purposes or with different inputs than what we considered in this paper. In software engineering, several techniques have been proposed to provide suggestions for certain fields of software systems (e.g., priority of bug reports in bug tracking systems [72] and elements of a domain model when using domain-model designing systems [14]). Suggestions are provided by analyzing the textual data (e.g., bug reports, requirement documents) with natural language processing techniques; instead, in our work we analyze the dependencies of categorical fields. In the field of information retrieval, crawlers automatically fill and submit web forms to crawl the data returned from databases [39, 47]; in the context of data crawling, automated form filling aims to automatically generate input values that can pass the field validation and retrieve more data, instead of helping users find the correct value they intend to fill. In the field of data mining, several approaches [22, 50, 69] fill data entry forms with the metadata extracted from data-rich text files. For example, these approaches can automatically use the metadata taken

from a resume text file to fill several fields of forms in different job search sites [69]. However, in this study we do not rely on data-rich text files to infer field values.

Table 5 presents an overview of related work and their differences with LAFF.

8 CONCLUSION

In this paper, we proposed LAFF, an approach to automatically suggest possible values of categorical fields in data entry forms, which are common user interface features in many software systems. Our approach utilizes Bayesian Networks to learn field dependencies from historical input instances. Moreover, LAFF relies on a clustering-based local modeling strategy to mine local field dependencies from partitions of historical input instances, to improve its learning ability. Furthermore, LAFF uses a heuristic-based endorser to ensure minimal accuracy for suggested values.

We evaluated LAFF by assessing its effectiveness and efficiency in form filling on two datasets, one of them proprietary from the banking domain. Evaluation results show that LAFF can provide a large number of accurate form filling suggestions, significantly outperforming state-of-the-art approaches in terms of Mean Reciprocal Rank (MRR). Further, LAFF takes at most 317 ms to provide a suggestion and is therefore applicable in practical data-entry scenarios.

As part of future work, we plan to conduct a study from the point of view of both users and developers, to analyze the effect of LAFF on reducing form filling time, input errors, and the cost of developing data entry forms. We also plan to assess the performance of LAFF under different deployment configurations and the accuracy of LAFF for different datasets and data entry forms.

Finally, we are also going to investigate methods to reduce the number of incorrect suggestions provided by LAFF when the number of filled fields used for prediction or the size of the training set is small.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. Financial support for this work was provided by the Alphonse Weicker Foundation.

REFERENCES

- [1] RB Aggarwal, Amit Dhawan, and Jay Shankar Kumar. 2007. Database-centric development of menus and graphic user interfaces. *Defence Science Journal* 57, 1 (2007), 133.
- [2] Pierre A Akiki, Arosha K Bandara, and Yijun Yu. 2016. Engineering adaptive model-driven user interfaces. *IEEE Transaction on Software Engineering* 42, 12 (2016), 1118–1147.
- [3] Stamatios-Aggelos N. Alexandropoulos, Sotiris B. Kotsiantis, and Michael N. Vrahatis. 2019. Data preprocessing in predictive data mining. *Knowledge Engineering Review* 34 (2019), e1.
- [4] Alnur Ali and Chris Meek. 2009. *Predictive Models of Form Filling*. Technical Report MSR-TR-2009-1. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/predictive-models-of-form-filling/>
- [5] American Medical News. 2005. Data entry is a top cause of medication errors. <https://amednews.com/article/20050124/profession/301249959/4/>.
- [6] Yuan An, Xiaohua Hu, and Il-Yeol Song. 2012. Learning to discover complex mappings from web forms to ontologies. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM, New York, NY, USA, 1253–1262.
- [7] Android API Reference. 2021. Android View Autofill. <https://developer.android.com/reference/kotlin/android/view/autofill/package-summary>.
- [8] Ankur Ankan and Abinash Panda. 2015. PgmPy: probabilistic graphical models using python. In *Proc. SCIPY'15*. SCIPY, Austin, Texas, USA, 6–11.
- [9] Samur Araujo, Qi Gao, Erwin Leonardi, and Geert-Jan Houben. 2010. Carbon: domain-independent automatic web form filling. In *Proc. ICWE'10 (LNCS, Vol. 6189)*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, Germany, 292–306.
- [10] Iman Avazpour, Teerat Pitakrat, Lars Grunske, and John Grundy. 2014. Dimensions and metrics for evaluating recommendation systems. In *Recommendation Systems in Software Engineering*. Springer, Berlin, Heidelberg, Germany,

245–273.

- [11] Bank for International Settlements. 2003. General Guide to Account Opening and Customer Identification. <https://www.bis.org/publ/bcbs85annex.htm>.
- [12] Tanya Barrett, Karen Clark, Robert Gevorgyan, Vyacheslav Gorenlenkov, Eugene Gribov, Ilene Karsch-Mizrachi, Michael Kimelman, Kim D Pruitt, Sergei Resenchuk, Tatiana Tatusova, et al. 2012. BioProject and BioSample databases at NCBI: facilitating capture and organization of metadata. *Nucleic acids research* 40, D1 (2012), D57–D63.
- [13] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. 1984. *Classification and regression trees*. CRC press, Boca Raton, Florida, USA.
- [14] Loli Burgueño, Robert Clarisó, Shuai Li, Sébastien Gérard, and Jordi Cabot. 2021. A NLP-based architecture for the autocompletion of partial domain models. In *Proc. CAiSE'21 (LNCS)*. Springer, Berlin, Heidelberg, Germany, 15 pages.
- [15] Valeria Cardellini, Michele Colajanni, and Philip S Yu. 2000. Geographic load balancing for scalable distributed web systems. In *Proc. MASCOOTS'00*. IEEE, San Francisco, CA, USA, 20–27.
- [16] Pablo Castells, Saúl Vargas, and Jun Wang. 2011. Novelty and diversity metrics for recommender systems: choice, discovery and relevance. In *Proc. DDR'11 - International Workshop on Diversity in Document Retrieval*. self-published, <http://www.dcs.gla.ac.uk/workshops/ddr2011/ddr2011.proceedings.pdf>, 29–36.
- [17] Kuang Chen, Harr Chen, Neil Conway, Joseph M Hellerstein, and Tapan S Parikh. 2011. Usher: Improving data quality with dynamic forms. *IEEE Transaction on Knowledge and Data Engineering* 23, 8 (2011), 1138–1153.
- [18] Tse-Hsun Chen, Weiye Shang, Jinqiu Yang, Ahmed E Hassan, Michael W Godfrey, Mohamed Nasser, and Parminder Flora. 2016. An empirical study on the practice of maintaining object-relational mapping code in Java systems. In *Proc. MSR'16*. ACM, New York, NY, USA, 165–176.
- [19] Takeshi Chusho, Katsuya Fujiwara, and Keiji Minamitani. 2002. Automatic filling in a form by an agent for web applications. In *Proc. APSEC'02*. IEEE, Berlin, Heidelberg, Germany, 239–247.
- [20] Andy Cockburn and Carl Gutwin. 2009. A predictive model of human performance with scrolling and hierarchical lists. *Human-Computer Interaction* 24, 3 (2009), 273–314.
- [21] Ofer Dekel, Ohad Shamir, and Lin Xiao. 2010. Learning to classify with missing and corrupted features. *Machine learning* 81, 2 (2010), 149–178.
- [22] Oscar Diaz, Itziar Otaduy, and Gorka Puente. 2013. User-driven automation of web form filling. In *Proc. ICWE'13 (LNCS, Vol. 7977)*. Springer, Berlin, Heidelberg, Germany, 171–185.
- [23] Khaled El Emam. 1999. Benchmarking Kappa: Interrater agreement in software process assessments. *Empirical Software Engineering* 4, 2 (1999), 113–133.
- [24] Wenfei Fan, Floris Geerts, and Xibe Jia. 2008. A revival of integrity constraints for data cleaning. *Proc. VLDB Endowment*'08 1, 2 (2008), 1522–1523.
- [25] Sergio Firmenich, Vincent Gaits, Silvia Gordillo, Gustavo Rossi, and Marco Winckler. 2012. Supporting users tasks with personal information management and web forms augmentation. In *Proc. ICWE2012 (LNCS, Vol. 7387)*. Springer, Oxford University Press, Berlin, Heidelberg, Germany, 268–282.
- [26] Susan Fowler and Victor Stanwick. 2004. *Web application design handbook: Best practices for web-based software*. Morgan Kaufmann, Amsterdam, Boston, USA.
- [27] Nir Friedman, Dan Geiger, and Moises Goldszmidt. 1997. Bayesian network classifiers. *Machine learning* 29, 2-3 (1997), 131–163.
- [28] Abdul Gafur. 2020. Updated tabular key and improved browser-based interactive key to species of Pratylenchus Filipjev (Nematoda: Pratylenchidae). *Biodiversitas Journal of Biological Diversity* 21, 8 (2020), 3780–3785.
- [29] José A Gámez, Juan L Mateo, and José M Puerta. 2011. Learning Bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood. *Data Mining and Knowledge Discovery* 22, 1-2 (2011), 106–148.
- [30] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proc. RecSys'10*. ACM, New York, NY, USA, 257–260.
- [31] Rafael S. Gonçalves, Martin J. O'Connor, Marcos Martínez-Romero, Attila L. Egyedi, Debra Willrett, John Graybeal, and Mark A. Musen. 2017. The CEDAR workbench: an ontology-assisted environment for authoring metadata that describe scientific experiments. In *Proc. ISWC'17 (LNCS, Vol. 10588)*. Springer International Publishing, Cham, 103–110.
- [32] Google. 2008. Chrome autofill forms. <https://support.google.com/chrome>. Accessed Feb 18, 2020.
- [33] Carl Gutwin and Andy Cockburn. 2006. Improving list revisitation with ListMaps. In *Proc. AVI'06*. ACM, New York, NY, USA, 396–403.
- [34] Melanie Hartmann and Max Muhlhauser. 2009. Context-aware form filling for web applications. In *Proc. ICSC'09*. IEEE, Berkeley, CA, USA, 221–228.
- [35] Bin He, Kevin Chen-Chuan Chang, and Jiawei Han. 2004. Discovering complex matchings across web query interfaces: a correlation mining approach. In *Proc. KDD'04*. ACM, New York, NY, USA, 148–157.
- [36] Carrie Heeter. 2000. Interactivity in the context of designed experiences. *J. of Interactive Advertising* 1, 1 (2000), 3–14.

- [37] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Transaction on Information Systems* 22, 1 (2004), 5–53.
- [38] L. A. Hermens and J. C. Shlimmer. 1994. A machine-learning apprentice for the completion of repetitive forms. *IEEE Expert* 9, 1 (1994), 28–33.
- [39] Inma Hernández, Carlos R Rivero, and David Ruiz. 2019. Deep web crawling: a survey. *World Wide Web* 22, 4 (2019), 1577–1610.
- [40] Jan Horsky, David R Kaufman, Michael I Oppenheim, and Vimla L Patel. 2003. A framework for analyzing the cognitive complexity of computer-assisted clinical ordering. *Journal of Biomedical Informatics* 36, 1-2 (2003), 4–22.
- [41] Zhexue Huang. 1998. Extensions to the k-means algorithm for clustering large data sets with categorical values. *Data Mining and Knowledge Discovery* 2, 3 (1998), 283–304.
- [42] Caroline Jarrett and Gerry Gaffney. 2009. *Forms that work: Designing Web forms for usability*. Morgan Kaufmann, Amsterdam, Boston, USA.
- [43] Emil Thorstensen Jensen, Martin Hansen, Evelyn Eika, and Frode Eika Sandnes. 2020. Country selection on web forms: a comparison of dropdown menus, radio buttons and text field with autocomplete. In *Proc. IMCOM'20*. IEEE, IEEE, Taiwan, China, 1–4.
- [44] Xiao-Yuan Jing, Fumin Qi, Fei Wu, and Baowen Xu. 2016. Missing data imputation based on low-rank recovery and semi-supervised regression for software effort estimation. In *Proc. ICSE'16*. IEEE, Austin, TX, USA, 607–618.
- [45] Chichang Jou. 2019. Schema Extraction for Deep Web Query Interfaces Using Heuristics Rules. *Information Systems Frontiers* 21, 1 (2019), 163–174.
- [46] Marius Kaminskis and Derek Bridge. 2016. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems* 7, 1 (2016), 1–42.
- [47] Gustavo Zanini Kantorski, Viviane Pereira Moreira, and Carlos Alberto Heuser. 2015. Automatic filling of hidden web forms: a survey. *ACM SIGMOD Record* 44, 1 (2015), 24–35.
- [48] Mozghan Karimi, Dietmar Jannach, and Michael Jugovac. 2018. News recommender systems—Survey and roads ahead. *Information Processing & Management* 54, 6 (2018), 1203–1227.
- [49] Reza Khajouei and MWM Jaspers. 2010. The impact of CPOE medication systems’ design aspects on usability, workflow and medication orders. *Methods of Information in Medicine* 49, 01 (2010), 03–19.
- [50] Trausti Kristjansson, Aron Culotta, Paul Viola, and Andrew McCallum. 2004. Interactive information extraction with constrained conditional random fields. In *Proc. AAAI'04*, Vol. 4. ACM, New York, NY, USA, 412–418.
- [51] Matevž Kunaver and Tomaž Požrl. 2017. Diversity in recommender systems—A survey. *Knowledge-based Systems* 123 (2017), 154–162.
- [52] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. 2019. Melu: Meta-learned user preference estimator for cold-start recommendation. In *Proc. SIGKDD'19*. ACM, New York, NY, USA, 1073–1082.
- [53] Marcos Martínez-Romero, Martin J O’Connor, Attila L Egyedi, Debra Willrett, Josef Hardi, John Graybeal, and Mark A Musen. 2019. Using association rule mining and ontologies to generate metadata recommendations from multiple biomedical databases. *Database J. Biol. Databases Curation* 2019 (2019), 25 pages.
- [54] S. McIntosh and Y. Kamei. 2017. Are fix-inducing changes a moving target? a longitudinal case study of just-in-time defect prediction. *IEEE Transaction on Software Engineering* 44, 5 (2017), 412–428.
- [55] Tim Menzies, Andrew Butcher, Andrian Marcus, Thomas Zimmermann, and David Cok. 2011. Local vs. global models for effort estimation and defect prediction. In *Proc. ASE'11*. IEEE, Lawrence, KS, USA, 343–351.
- [56] Microsoft. 2013. Change the default tab order for controls on a form. <https://support.microsoft.com/en-us/office/change-the-default-tab-order-for-controls-on-a-form-03d1599a-debf-4b66-a95b-e3e744210afe>.
- [57] Kıvanç Muşlu, Yuriy Brun, and Alexandra Meliou. 2015. Preventing data errors with continuous testing. In *Proc. ISSTA'15*. ACM, New York, NY, USA, 373–384.
- [58] John-Paul Ore, Sebastian Elbaum, Carrick Detweiler, and Lambros Karkazis. 2018. Assessing the type annotation burden. In *Proc. ASE'18*. ACM, New York, NY, USA, 190–201.
- [59] Spencer Pearson, José Campos, René Just, Gordon Fraser, Rui Abreu, Michael D Ernst, Deric Pang, and Benjamin Keller. 2017. Evaluating and improving fault localization. In *Proc. ICSE'17*. IEEE, Buenos Aires, Argentina, 609–620.
- [60] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [61] Siyu Qian, Esther Munyisia, David Reid, David Hailey, Jade Pados, and Ping Yu. 2020. Trend in data errors after the implementation of an electronic medical record system: A longitudinal study in an Australian regional Drug and Alcohol Service. *International Journal of Medical Informatics* 144 (2020), 104292.
- [62] Adrian E Raftery. 1995. Bayesian model selection in social research. *Sociological Methodology* 25 (1995), 111–163.

- [63] Martin Robillard, Robert Walker, and Thomas Zimmermann. 2009. Recommendation systems for software engineering. *IEEE Software* 27, 4 (2009), 80–86.
- [64] Enrico Rukzio, Chie Noda, Alexander De Luca, John Hamard, and Fatih Coskun. 2008. Automatic form filling on mobile devices. *Pervasive and Mobile Computing* 4, 2 (2008), 161–181.
- [65] Marton Sakal and Lazar Rakovic. 2012. Errors in building and using electronic tables: Financial consequences and minimisation techniques. *Strategic Management* 17, 3 (2012), 29–35.
- [66] Amr Rekaby Salama, Ozge Alaçam, and Wolfgang Menzel. 2018. Text completion using a context-integrating dependency parser. In *Proc. RepL4NLP'18*. ACL, Melbourne, Australia, 41–49.
- [67] Gunnar Schröder, Maik Thiele, and Wolfgang Lehner. 2011. Setting goals and choosing metrics for recommender system evaluations. In *UCERSTI2 workshop at the 5th ACM conference on recommender systems*, Vol. 23. ACM, New York, NY, USA, 53.
- [68] Andrew Sears and Ying Zha. 2003. Data entry for mobile devices using soft keyboards: Understanding the effects of keyboard size and user tasks. *J. of Human-Computer Interaction* 16, 2 (2003), 163–184.
- [69] Guilherme A Toda, Eli Cortez, Altigran S da Silva, and Edleno de Moura. 2010. A probabilistic approach for automatically filling form-based web interfaces. *Proc. of the VLDB Endowment* 4, 3 (2010), 151–160.
- [70] Clare L Tolley, Niamh E Forde, Katherine L Coffey, Dean F Sittig, Joan S Ash, Andrew K Husband, David W Bates, and Sarah P Slight. 2018. Factors contributing to medication errors made when using computerized order entry in pediatrics: a systematic review. *Journal of the American Medical Informatics Association* 25, 5 (2018), 575–584.
- [71] Luigi Troiano, Cosimo Birtolo, and Roberto Armenise. 2017. Modeling and predicting the user next input by Bayesian reasoning. *Soft Computing* 21, 6 (2017), 1583–1600.
- [72] Qasim Umer, Hui Liu, and Inam Illahi. 2019. CNN-based automatic prioritization of bug reports. *IEEE Transaction on Reliability* 69, 4 (2019), 1341–1354.
- [73] Antal Van Den Bosch and Toine Bogers. 2008. Efficient context-sensitive word completion for mobile devices. In *Proc. MobileHCI'08*. ACM, New York, NY, USA, 465–470.
- [74] W3C School. 2021. HTML <input> autocomplete Attribute. https://www.w3schools.com/tags/att_input_autocomplete.asp.
- [75] W3CSchools. 2017. HTML DOM input text object. https://www.w3schools.com/jsref/dom_obj_text.asp.
- [76] Shaohua Wang, Ying Zou, Iman Keivanloo, Bipin Upadhyaya, Joanna Ng, and Tinny Ng. 2014. Automatic reuse of user inputs to services among end-users in service composition. *IEEE Transaction on Services Computing* 8, 3 (2014), 343–355.
- [77] Shaohua Wang, Ying Zou, Joanna Ng, and Tinny Ng. 2017. Context-aware service input ranking by learning from historical information. *IEEE Transaction on Services Computing* 14, 1 (2017), 97–110.
- [78] Johanna I Westbrook, Melissa T Baysari, Ling Li, Rosemary Burke, Katrina L Richardson, and Richard O Day. 2013. The safety of electronic prescribing: manifestations, mechanisms, and rates of system-related errors associated with two commercial systems in hospitals. *Journal of the American Medical Informatics Association* 20, 6 (2013), 1159–1167.
- [79] Marco Winckler, Vicent Gaits, Dong-Bach Vo, Firmenich Sergio, and Gustavo Rossi. 2011. An approach and tool support for assisting users to fill-in web forms with personal information. In *Proc. SIGDOC'11*. ACM, New York, NY, USA, 195–202.
- [80] Xinli Yang, David Lo, Xin Xia, Lingfeng Bao, and Jianling Sun. 2016. Combining word embedding with information retrieval to recommend similar bug reports. In *Proc. ISSRE'16*. IEEE, Ottawa, ON, Canada, 127–137.
- [81] Xin Ye, Razvan Bunescu, and Chang Liu. 2014. Learning to rank relevant files for bug reports using domain knowledge. In *Proc. FSE'14*. ACM, New York, NY, USA, 689–699.
- [82] Mingrui Ray Zhang, Shumin Zhai, and Jacob O. Wobbrock. 2019. Text entry throughput: towards unifying speed and accuracy in a single performance metric. In *Proc. CHI'19*. ACM, New York, NY, USA, 1–13.
- [83] Zhi-Hua Zhou. 2021. Ensemble learning. In *Machine Learning*. Springer, Singapore, 181–210.