



# Learning and Deducing Temporal Orders

Wenfei Fan  
Shenzhen Institute of  
Computing Sciences, China  
University of Edinburgh  
United Kingdom  
Beihang University, China  
wenfei@inf.ed.ac.uk

Resul Tugay  
University of Edinburgh  
United Kingdom  
resul.tugay@ed.ac.uk

Yaoshu Wang\*  
Min Xie  
Shenzhen Institute of  
Computing Sciences, China  
yaoshuw@sics.ac.cn  
xiemin@sics.ac.cn

Muhammad Asif Ali  
King Abdullah University  
of Science and Technology  
Kingdom of Saudi Arabia  
muhammadasif.ali  
@kaust.edu.sa

## ABSTRACT

This paper studies how to determine temporal orders on attribute values in a set of tuples that pertain to the same entity, in the absence of complete timestamps. We propose a creator-critic framework to learn and deduce temporal orders by combining deep learning and rule-based deduction, referred to as GATE (Get the *LATE*st). The creator of GATE trains a ranking model via deep learning, to learn temporal orders and rank attribute values based on correlations among the attributes. The critic then validates the temporal orders learned and deduces more ranked pairs by chasing the data with currency constraints; it also provides augmented training data as feedback for the creator to improve the ranking in the next round. The process proceeds until the temporal order obtained becomes stable. Using real-life and synthetic datasets, we show that GATE is able to determine temporal orders with  $F$ -measure above 80%, improving deep learning by 7.8% and rule-based methods by 34.4%.

## PVLDB Reference Format:

Wenfei Fan, Resul Tugay, Yaoshu Wang, Min Xie, and Muhammad Asif Ali. Learning and Deducing Temporal Orders. PVLDB, 16(8): 1944 - 1957, 2023. doi:10.14778/3594512.3594524

## PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/resultugay/GATE>.

## 1 INTRODUCTION

Real-life data keeps changing. As reported by Royal Mail, on average, “9590 households move, 1496 people marry, 810 people divorce, 2011 people retire and 1500 people die” each day in the UK [60]. It is estimated that inaccurate customer data costs organizations 6% of their annual revenues [60]. Outdated data incurs damage not only to Royal Mail. When the data at a search engine is out of date, a restaurant search may return a business that had closed three years ago. When the data about the condition of infrastructure assets is obsolete, it may delay the maintenance of equipment and cause outage. Moreover, data-driven decisions based on outdated data can be worse than making decisions with no data [49]. Indeed, “as a healthcare, retail, or financial services business you cannot afford to

make decisions based on yesterday’s data” [25]. Unfortunately, “82% of companies are making decisions based on stale information” [8].

These highlight the need for determining the currency of data, *i.e.*, how up-to-date the information is. This is, however, highly nontrivial. Consider a set of tuples that pertain to the same entity. Their attribute values may become obsolete and inaccurate over the time. Worse yet, only *partial* reliable timestamps might be available, where a timestamp is *reliable* if it is *precise*, *correct* and moreover, it indicates that at the time, the values are *correct* and *up-to-date*. Apart from mechanical reasons (malicious attacks or hardware failures), the logical reasons below are the major temporal issues in data quality [5], and account for the absence of reliable timestamps.

(1) *Missing timestamps*. Timestamps may simply not be recorded, *e.g.*, in an e-health database [40], only 16 out of 26 relations are timestamped. Even when a relation has timestamps, it may not be complete, *e.g.*, 25.36% missing in [53], up to 82.28% in [40].

(2) *Imprecise timestamps*. Timestamps may be too coarse, leading to unreliable ordering. An e-form may be submitted multiple times to an OA system by employees during a day. If the forms are recorded in datestamps, it is not clear which form (all on the same day) is the latest. Similar problems are often encountered in hospital data, where only dates are recorded [5]. Another example concerns inconsistent granularity of timestamps (*e.g.*, minutes vs. days [5, 40]). If two values have timestamps “12-8-2021” and “12-8-2021 20:41”, it is not clear which one is more up-to-date. As reported in [53], 90.41% of appointment records have imprecise timestamps.

(3) *Incorrect timestamps*. Many factors can lead to incorrect timestamps. Taking medical data [5] as an example, an X-ray machine has many asynchronous modules, each of which has a local clock and a local buffer. There can be a discrepancy between when a value is actually updated and when it is recorded since the value is first queued in the buffer before it is recorded. Moreover, 36.96% of appointments have the overlapping issue (*i.e.*, the next appointment in a particular room appears to have started before the current one has ended) [53], indicating incorrect timestamps.

As remarked earlier by Royal Mail [60], customer data changes frequently. To prevent the data from being outdated, the sales department may contact the customers and get regular updates on some *critical* information (*e.g.*, email and phone). Due to the impatience of customers and high cost (*e.g.*, man power) of contact, the rest (*e.g.*, jobs, marital status) may not be frequently updated and may gradually become obsolete, no longer having a reliable timestamp. Add to this the complication that data values are often copied or imported from other sources [21, 22], and even in the

\*Corresponding author

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 16, No. 8 ISSN 2150-8097.  
doi:10.14778/3594512.3594524

	FN	LN	sex	address	status	job	kids	SZ
$t_1, e_1$ :	Mary	Goldsmith	F	19 xin st	single	n/a	-	5
$t_2, e_1$ :	Mary	Taylor	F	6 gold plaza	married	journalist	1	6
$t_3, e_1$ :	Mary	Taylor	F	19 mall st	espoused	assoc editor	2	7
$t_4, e_1$ :	Mary	Taylor	F	7 ave	divorced	chief editor	2	7
$t_5, e_1$ :	Mary	Taylor	F	7 ave	detached	chief editor	2	7
$t_6, e_1$ :	Mary	Goldsmith	F	6 const. ave	wed	producer	3	7

**Figure 1: Customer records dataset**

same database, values may come from multiple relations (*e.g.*, by join), where no uniform scheme of timestamps is granted, *e.g.*, some values are more frequently timestamped than the others [40]. This calls for an attribute-level time-stamping scheme (*i.e.*, each value has a timestamp). It subsumes the tuple-level scheme (*i.e.*, all values in the same tuple have the same timestamp) as a special case.

No matter how desirable, the percentage of reliable timestamps is far lower than expected. How can we determine the temporal orders on the attribute values, *i.e.*, for tuples  $t_1$  and  $t_2$ , whether the value in the  $A$ -attribute of  $t_1$  is more current than the value in the  $A$ -attribute of  $t_2$ , denoted by  $t_2 <_A t_1$ , in the absence of complete timestamps?

**Example 1:** Consider customer records  $t_1$ - $t_6$  shown in Figure 1, which have been identified to refer to the same person Mary. Each tuple  $t_i$  has attributes FN (first name), LN (last name), sex, address, marital status, job, kids and SZ (shoe size). Some attribute values of these tuples have become stale since Mary’s data changes over the years, *e.g.*, her job, address and last name have changed 4 times, five times and twice, respectively. Only some attribute values might be associated with reliable timestamps, *e.g.*, the timestamp of  $t_5$ [job] and  $t_6$ [job] are 2016 and 2019 (not shown), respectively, indicating that at that time, the values are up-to-date. In the absence of complete timestamps, it is hard to know whether  $t_2 <_{LN} t_6$ , *i.e.*, whether the value of  $t_2$ [LN] is more up-to-date than  $t_6$ [LN]? Moreover, what Mary’s current job title is, *e.g.*, whether  $t_i <_{job} t_6$  for  $i \in [1, 4]$ ? □

One may want to approach this problem by training a ranking model that sorts objects according to their degrees of relevance or importance [51]. The state-of-the-art systems in this regard employ deep learning [7, 36, 57, 63] or reinforcement learning [37, 68], and have been used in search engine [11] and machine translation [24, 70]. By means of ranking models one can learn temporal orders and decide whether  $t_1 <_A t_2$  for *all* tuples  $t_1, t_2$  and attribute  $A$ . However, it is hard to justify whether the ranking conforms to the temporal orders in the real world. For data-driven decision making, we need to ensure that the learned orders are reliable. Moreover, these approaches cannot explain the ranking of objects that follow a complex interlinked structure (*e.g.*, address in Figure 1).

Another approach is to employ logic rules, *e.g.*, currency constraints (CCs) [19, 23, 28, 30, 44, 48, 66]. These rules help us deduce temporal orders. Taking Example 1 as an example, the shoe sizes of the same person typically monotonically increases (before 20 years old), and the address of a person may be associated with the marital status. As will be seen in Section 2, using CCs one can deduce that  $t_2 <_{SZ} t_3$  and  $t_1 <_{address} t_2$ . However, it is hard to find enough rules to deduce relative orders on *each and every* pair of values. For instance, we cannot find rules to decide whether  $t_2 <_{LN} t_6$ . When testing existing rule-based methods on a dataset with 5% initial timestamps, even the best of them can only deduce 16.3% temporal orders (Section 6), leaving the remaining 78.7% undecided. Besides, it is hard to generalize rules to handle lexically different

but semantically similar values, since tuples might be extracted from different sources (*e.g.*, status in Figure 1: married vs. wed).

Neither deep/ML learning nor logic rules work well when used separately. A natural question is whether it is possible to combine ML models and logic rules in a uniform framework, such that we can learn temporal orders, and use the rules to validate the ranking and improve the learning? How well can this framework improve the accuracy of the ML models and logic rules when being used alone?

**Contributions & organization.** This paper makes a first effort to answer these questions, all in the affirmative.

(1) *Temporal orders* (Section 2). We define the notion of temporal orders  $t_1 <_A t_2$  and  $t_1 \leq_A t_2$  on attributes, and formulate the problem for determining temporal orders. We also review the currency constraints (CCs) of [30], for deducing temporal orders by our framework. We show that CCs can specify interesting temporal properties, *e.g.*, the monotonicity, comonotonicity and transitivity.

(2) *GATE* (Section 3). We propose GATE (Get the  $LATEst$ ), a creator-critic framework to determine temporal orders by combining deep learning and logic deduction. GATE iteratively invokes a creator to rank the temporal orders on attribute values, followed by the critic that validates the ranking of the creator and deduces more ranked pairs via discovered CCs. The critic also produces augmented training data for the creator to improve its ranking in the next round. This process proceeds for the creator and critic to mutually enhance each other, until the temporal order cannot be further improved.

(3) *Creator* (Section 4). We propose a deep learning model underlying the creator of GATE, to learn temporal orders on attribute values. It departs from previous ranking models in that it learns the temporal orders based on contexts (*i.e.*, attribute correlations) and calculates the confidence of the ranking, by employing chronological embeddings and adaptive pairwise ranking strategies.

(4) *Critic* (Section 5). The critic complements GATE with discovered rules (CCs). We show how it justifies the ranked pairs learned by the creator, and (incrementally) deduces latent temporal orders with the chase [62] using CCs. The chase has the Church-Rosser property (cf. [2]), *i.e.*, it guarantees to converge at the same result no matter in what orders the CCs are applied. Moreover, the critic augments the training data for the creator to improve its model in the next round.

(5) *Experimental study* (Section 6). Using real-life and synthetic data, we find that (a) the  $F$ -measure of GATE on dataset Career [41] is 0.866, versus 0.41 and 0.39 by rule-based UncertainRule [44] and Improve3C [18] (resp. 0.54 and 0.53 by ML-based RANK<sub>Bert</sub> [56] and Ditto<sub>Rank</sub> [47]). (b) On average, GATE is 34.4% and 7.8% more accurate than the critic and the creator, respectively, verifying the effectiveness of combining deep learning and logic rules. (c) GATE is feasible in practice; it only takes 7 rounds to terminate on a real-life dataset of 1,983,698 tuples, with a single machine.

**Related work.** We categorize the related work as follows.

*Rule-based methods.* Currency constraints (CCs) were first studied in [29, 30] by employing partial currency orders and later extended in [27, 28] for conflict resolution, by considering both CCs and conditional functional dependencies (CFDs) [26]. A class of currency repairing rules (CRRs) was proposed in [43], which combines logic rules and statistics. A two-step approach was developed in [17] to

determine the currency of dynamic data, by means of a topological graph. A class of uncertain currency rules was supported by UncertainRule [44] that considers both temporal orders and data certainty. Improve3C [18] and Imp3C [19] are 4-step data cleaning frameworks, including data consistency, completeness and currency, which use a metric to repair noisy data using CFDs [26] and currency orders. There has also been work [23] on parallel incremental updating algorithms by employing traditional currency rules.

This work differs from the prior work. To the best of our knowledge, we make the first effort to combine deep learning and rules for timeliness, for the two to enhance each other, while the prior work at most extends rules with statistic. We propose a model for inferring temporal orders and use logic deduction to derive more.

**ML models.** There have also been efforts on inferring temporal orders by ML models [6, 10, 35, 54, 55, 64]. A related topic is to incorporate temporal information into knowledge graphs, e.g., for temporal link predictions [20, 61] and reasoning [16, 65, 69]. These methods often embed temporal information in ML models, for inference varying over time. Learning temporal orders has been modeled as a learning-to-rank problem (pointwise, pairwise and listwise), where global orders are learned for currency attributes. Temporal problems have also been studied in, e.g., search and recommendation [31, 67], IR [50] and NLP [47, 56]. In particular, ranking in NLP (e.g., in question-answering tasks) can be handled by pre-trained language models with cross-entropy loss or by a ranking function with a binary classifier as the comparison operator; we adopted a baseline from each kind (i.e., RANK<sub>Bert</sub> and Ditto<sub>Rank</sub>) in Section 6.

In contrast, we propose a creator-critic framework that not only learns temporal orders via deep learning, but also adopts rules to deduce new ones and help ML models learn better. The prior models can be embedded into our framework.

**Entity resolution (ER).** Temporal clustering methods were proposed in [46] based on a concept of time decay, to improve ER with timestamps. [12] dynamically assesses and adjusts similarities among records based on their attributes and temporal differences. [42] estimates the probability that a value changes, and shows how to link tuples in a correct time period. To cope with incorrect timestamps, [66] adopts matching dependencies (MDs) and currency constraints to capture the attributes that were changed over time.

This work focuses on deducing temporal orders of each entity, while the prior work mainly adopts temporal information for ER.

## 2 DETERMINING TEMPORAL ORDERS

In this section we first formulate temporal orders and the problem for determining temporal orders (Section 2.1). We then review currency constraints (CCs) of [30] and show that such rules are able to express monotonicity, comonotonicity and transitivity (Section 2.2).

### 2.1 The Problem

Consider a relation schema  $R = (EID, A_1, \dots, A_n)$ , where  $A_i$  is an attribute ( $i \in [1, n]$ ), and EID is an entity id as introduced by Codd [14]. For a tuple  $t$  of schema  $R$ , we denote by  $t[A]$  the value in the  $A$ -attribute of  $t$ , and by  $t[EID]$  the entity represented by tuple  $t$ . We assume that each tuple is associated with a tuple id, denoted by tid. We refer to a relation  $D$  of  $R$  as a *normal instance*  $D$  of  $R$ .

**Temporal orders and temporal instances.** An *entity instance* is  $(D_e, T_e)$ , where (a)  $D_e$  is a normal instance of schema  $R$  such that all tuples  $t_1$  and  $t_2$  in  $D_e$  refer to the same real-life entity  $e$  and hence,  $t_1[EID] = t_2[EID]$ ; and (b)  $T_e$  is a partial function that associates a timestamp  $T_e(t[A])$  with the  $A$ -attribute of a tuple  $t$  in  $D_e$ . We refer to  $(D_e, T_e)$  as an *entity instance pertaining to  $e$* .

Here the timestamp indicates that at the time  $T_e(t[A])$ , the  $A$ -attribute value of tuple  $t$  is correct and up-to-date; it does not necessarily refer to the time when  $t[A]$  was created or last updated. If  $T_e(t[A])$  is undefined, a reliable timestamp is not available for  $t[A]$ .

Intuitively, an entity instance extends a normal instance with available timestamps. Its tuples may be extracted from a variety of data sources, and are identified to refer to the same entity  $e$  via entity resolution. In the same tuple  $t$ ,  $t[A]$  and  $t[B]$  may bear different timestamps (or even no timestamp) for different  $A$  and  $B$ . Note that we do not assume a timestamp for the entire tuple, since we often find only parts of a tuple to be correct and up-to-date.

**Temporal orders.** A *temporal order* on attribute  $A$  of  $D_e$  is a partial order  $\leq_A$  such that for all tuples  $t_1$  and  $t_2 \in D_e$ ,  $t_2 \leq_A t_1$  if the value in  $t_1[A]$  is at least as current as  $t_2[A]$ . We also use a strict partial order  $<_A t_1$  if  $t_1[A]$  is more current than  $t_2[A]$ . To simplify the discussion we focus on  $\leq_A$  in the sequel;  $<_A$  is handled analogously.

In particular, if  $T_e(t_1[A])$  and  $T_e(t_2[A])$  are both defined and if  $T_e(t_2[A]) \leq T_e(t_1[A])$ , i.e., when timestamp  $T_e(t_1[A])$  is no earlier than  $T_e(t_2[A])$ , then  $t_2 \leq_A t_1$ , i.e.,  $t_1[A]$  is confirmed at a later timestamp and is thus considered at least as current as  $t_2[A]$ .

Temporal order  $\leq_A$  is represented as a set of tuple pairs such that  $(t_2[\text{tid}], t_1[\text{tid}]) \in \leq_A$  iff  $t_2 \leq_A t_1$ . We write  $(t_2[\text{tid}], t_1[\text{tid}])$  as  $(t_2, t_1)$  if it is clear in the context. Note that the same value may bear different timeliness in different tuples, e.g., Mary's marital status changed from married ( $t_2$ ) to divorced ( $t_4$ ) to married ( $t_7$ , not shown in Figure 1). While  $t_2[\text{status}] = t_7[\text{status}]$ ,  $t_2 <_{\text{status}} t_7$ . Here  $t_2 <_{\text{status}} t_7$  ranks the timeliness of the status-attributes of tuples  $t_2$  and  $t_7$ , not values (married vs. married) detached from the tuples.

We say that a temporal order  $\leq_A^1$  extends  $\leq_A^2$ , written as  $\leq_A^2 \subseteq \leq_A^1$ , if for all tuples  $t_1, t_2$  in  $D_e$ , if  $t_2 \leq_A^2 t_1$  is defined, then so is  $t_2 \leq_A^1 t_1$ . That is,  $\leq_A^1$  includes all tuple pairs in  $\leq_A^2$  and possibly more.

**Temporal instances.** A *temporal instance*  $D_t$  of  $R$  is given as  $(D, \leq_{A_1}, \dots, \leq_{A_n}, T)$ , where each  $\leq_{A_i}$  is a temporal order on  $A_i$  ( $i \in [1, n]$ ),  $D = \bigcup_{i \in [k]} D_{e_i}$ ,  $T = \bigcup_{i \in [k]} T_{e_i}$ , and for all  $i \in [1, k]$ ,  $(D_{e_i}, T_{e_i})$  is an entity instance of  $R$ . Here tuples  $t_1$  and  $t_2$  in  $D$  are *compatible* under  $\leq_A$  if they pertain to the same entity, i.e.,  $t_1[EID] = t_2[EID]$ .

Intuitively,  $D_t$  is a collection of entity instances, such that each  $(D_{e_i}, T_{e_i})$  pertains to the same entity  $e_i$ . We do not rank the currency of tuples if they refer to different entities. A temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  is said to *extend* another temporal instance  $D'_t = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T)$  if for all  $i \in [1, n]$ ,  $\leq_{A_i}$  extends  $\leq'_{A_i}$ .

**Problem statement.** We study the problem for *determining the temporal orders* of a temporal instance, stated as follows.

- *Input:* A temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$ .
- *Output:* An extended temporal instance  $D'_t = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T)$  such that for all  $i \in [1, n]$ , (a)  $\leq'_{A_i}$  extends  $\leq_{A_i}$  and (b)  $\leq'_{A_i}$  is a total order on all compatible  $t_1$  and  $t_2$  with  $t_1[EID] = t_2[EID]$ .

Intuitively, our goal is to extend  $\leq_{A_i}$  such that for all tuples  $t_1$  and  $t_2$  in  $D$ , if  $t_1[\text{EID}] = t_2[\text{EID}]$ , we can decide which of  $t_1[A_i]$  and  $t_2[A_i]$  is more up-to-date. As a consequence, we can deduce the latest value for all attributes. Note that we define a total order on *each* attribute, not a global order on tuples. The total orders on different attributes can be different. This said, we can use the temporal orders learned on one attribute to help the deduction on other attributes, via correlation expressed as currency constraints (see below).

## 2.2 Currency Constraints

We next review the class of currency constraints proposed in [30]. The predicates over a relation schema  $R$  are defined as:

$$p ::= t[A] \oplus c \mid t_1[A] \oplus t_2[A] \mid t_1 \leq_A t_2,$$

where  $t, t_1, t_2$  are tuple variables denoting tuples of  $R$ ,  $A$  is an attribute of  $R$ ,  $c$  is a constant,  $\oplus$  is an operator from  $\{=, \neq, >, \geq, <, \leq\}$ ;  $t[A] \oplus c$  and  $t_1[A] \oplus t_2[A]$  are defined on attribute values, while  $t_1 \leq_A t_2$  compares the timeliness of  $t_1[A]$  and  $t_2[A]$ . In particular,  $t_1[\text{EID}] = t_2[\text{EID}]$  says that  $t_1$  and  $t_2$  refer to the same entity.

A *currency constraint* (CC) over schema  $R$  is defined as follows:

$$\varphi = X \rightarrow p_0,$$

where  $X$  is a conjunction of predicates over  $R$  with tuple variables  $t_1, \dots, t_m$ , and  $p_0$  has the form  $t_u \leq_{A_i} t_v$  for  $u, v \in [1, m]$ . We refer to  $X$  as the *precondition* of  $\varphi$ , and  $p_0$  as the *consequence* of  $\varphi$ .

As defined in [30], a CC can be equivalently expressed as a universal first-order logic sentence of the following form:

$$\varphi = \forall t_1, \dots, t_m \left( \bigwedge_{j \in [1, m]} (t_1[\text{EID}] = t_j[\text{EID}]) \wedge X \rightarrow t_u \leq_A t_v \right).$$

**Semantics.** Currency constraints are defined over temporal instances  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  of  $R$ . A *valuation* of tuple variables of a CC  $\varphi$  in  $D_t$ , or simply a *valuation of  $\varphi$* , is a mapping  $h$  that instantiates variables  $t_1, \dots, t_m$  with tuples in  $D$  that refer to the same real-world entity, as required by  $t_1[\text{EID}] = t_j[\text{EID}]$ .

A valuation  $h$  *satisfies* a predicate  $p$  over  $R$ , written as  $h \models p$ , if the following is satisfied: (1) if  $p$  is  $t[A] \oplus c$  or  $t_1[A] \oplus t_2[A]$ , then it is interpreted as in tuple relational calculus following the standard semantics of first-order logic [2]; and (2) if  $p$  is  $t_1 \leq_A t_2$ , then  $t_2[A]$  is at least as current as  $t_1[A]$  i.e.,  $(t_1, t_2)$  is in  $\leq_A$ .

For a conjunction  $X$  of predicates, we write  $h \models X$  if  $h \models p$  for all  $p$  in  $X$ . A temporal instance  $D_t$  *satisfies* CC  $\varphi$ , denoted by  $D_t \models \varphi$ , if for all valuations  $h$  of  $X \rightarrow p_0$  in  $D_t$ , if  $h \models X$  then  $h \models p_0$ .

**Properties.** Currency constraints are able to specify interesting temporal properties. Below we exemplify some properties.

**Monotonicity.** A temporal order  $\leq_A$  over relation schema  $R$  is *monotonic* if for any tuples  $t_1$  and  $t_2$  of  $R$  that refer to the same entity, if  $t_1[A] \leq t_2[A]$  then  $t_1 \leq_A t_2$ . For instance, consider the SZ (shoe size) attribute of the customer relation of Example 1. Then  $\leq_{\text{SZ}}$  is monotonic. It can be expressed as the following CC:

$$\varphi_1 = t_1[\text{SZ}] \leq t_2[\text{SZ}] \rightarrow t_1 \leq_{\text{SZ}} t_2.$$

As another example, marital status only changes from single to married, not the other way around [9]. This is expressed as CC:

$$\varphi_2 = t_1[\text{status}] = \text{"single"} \wedge t_2[\text{status}] = \text{"married"} \rightarrow t_1 \leq_{\text{status}} t_2.$$

With slight abuse of terminologies by considering timestamps as an "attribute" associated with attribute  $A$ , we have:

$$\varphi_3 = T_e(t_1[A]) \leq T_e(t_2[A]) \rightarrow t_1 \leq_A t_2,$$

since  $t_2[A]$  is confirmed at a later timestamp.

**Comonotonicity.** For attributes  $A$  and  $B$  of  $R$ , we say  $\leq_A$  and  $\leq_B$  are *comonotonic* in a temporal instance  $D_t$  of  $R$  if for all tuples  $t_1$  and  $t_2$  in  $D_t$  that refer to the same entity,  $t_1 \leq_A t_2$  if and only if  $t_1 \leq_B t_2$ .

Intuitively,  $\leq_A$  and  $\leq_B$  are comonotonic if the two are correlated such that when one is updated, the other will also change. For instance,  $\leq_{\text{status}}$  and  $\leq_{\text{address}}$  are often comonotonic: when the marital status of a person changes from single to married, this person may move to a larger house. This can be expressed as a CC:

$$\varphi_4 = t_1 \leq_{\text{status}} t_2 \rightarrow t_1 \leq_{\text{address}} t_2.$$

**Transitivity.** Transitivity can also be expressed for any attribute  $A$ :

$$\varphi_5 = t_1 \leq_A t_2 \wedge t_2 \leq_A t_3 \rightarrow t_1 \leq_A t_3.$$

**Correlating different attributes.** One can correlate multiple different attributes to capture implicit temporal ordering. For example, marital status may change from "married" to "divorced" and further from "divorced" to "married" again. To deduce an order on  $t[\text{status}]$ , we can use the number of kids as an additional condition:

$$\varphi_6 = t_1[\text{status}] = \text{"married"} \wedge t_2[\text{status}] = \text{"divorced"} \wedge t_1[\text{kids}] < t_2[\text{kids}] \rightarrow t_1 \leq_{\text{status}} t_2.$$

**Deduction.** Making use of CCs, we can deduce certain temporal orders, e.g., from  $\varphi_1$ - $\varphi_6$  and Figure 1, we can deduce the following:

- $t_1 \leq_{\text{status}} t_2$  by  $\varphi_2$ ,  $t_2 \leq_{\text{status}} t_4$  by  $\varphi_6$  and  $t_1 \leq_{\text{status}} t_4$  by  $\varphi_5$ ;
- $t_1 \leq_{\text{address}} t_2$  by  $\varphi_4$ ; hence  $t_2[\text{address}]$  is more current for Mary;
- $t_1 \leq_{\text{SZ}} t_2 \leq_{\text{SZ}} t_3$  by  $\varphi_1$ ; hence Mary's current shoe size is 7.

However, we cannot determine whether  $t_2 \leq_{\text{LN}} t_6$  or  $t_6 \leq_{\text{LN}} t_2$  by deduction with the currency constraints  $\varphi_1$ - $\varphi_6$ .

**Discovery of CCs.** As noted in [30], CCs are a special case of denial constraints (DCs) [3] extended with temporal orders  $\leq_A$  and  $<_A$ . Algorithms are in place for discovering DCs, e.g., [4, 13, 52, 58]. We can readily extend these algorithms for discovering CCs (see [1]).

## 3 GATE: A CREATOR-CRITIC SYSTEM

In this section, we propose a creator-critic framework for determining temporal orders, and develop system GATE to implement it. A unique feature of GATE is its combined use of deep learning and logic deduction. Below we start with the architecture of GATE, and then present its overall workflow, with termination guarantee.

**Architecture.** The ultimate goal of GATE is to obtain a total order  $\leq_A$  for each attribute  $A$ . As shown in Figure 2, GATE first discovers a set  $\Sigma$  of CCs on  $D_t$  *offline* for performing logic deduction. Then it takes a temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  as input, and learns and deduces more temporally ranked pairs for  $D_t$  *online*. For simplicity, we assume *w.l.o.g.* that  $D$  consists of a single entity instance  $D_e$ , i.e., all tuples in  $D$  pertain to the same entity  $e$  and thus their attributes can be pairwise compared; the methods of this paper can be readily extended to  $D_t$  with multiple entity instances.

More specifically, the learning and deducing process in GATE *iteratively* executes two phases, namely, creator and critic, as follows.

(1) *Creator* (Sections 4). In this phase, GATE (incrementally) trains a ranking model  $\mathcal{M}_{\text{rank}}$  via deep learning. By taking  $D_t$  and the augmented training data  $D_{\text{aug}}$  (see its definition shortly) from the critic as input, it predicts new orders for extending each  $\leq_A$  of  $D_t$ .

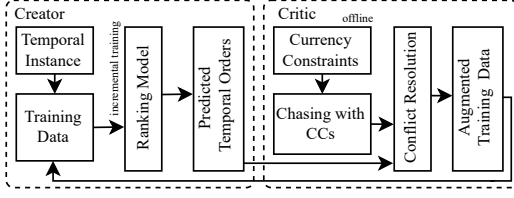


Figure 2: Architecture of GATE

Our model has three features: (a) For each tuple  $t$  in  $D$ , the embedding for  $t[A]$  is created using both  $t[A]$  and other correlated values in  $t$ , so that  $t[A]$  can be ranked comprehensively. (b) The attribute embeddings [15] are arranged to preserve chronological orders. (c) We adopt an *attribute-centric adaptive pairwise ranking* strategy in  $\mathcal{M}_{\text{rank}}$ , so that the ranking result can be justified semantically. Moreover, the temporal instance  $D_t$  will also be extended based on  $D_{\text{aug}}$ .

Given an attribute  $A$ , we associate each  $(t_1, t_2) \in \leq_A$  with a *confidence*, denoted by  $\text{conf}(t_1 \leq_A t_2)$ , indicating how likely  $t_1 \leq_A t_2$  holds. If  $t_2[A]$  has a later timestamp than  $t_1[A]$ , then  $\text{conf}(t_1 \leq_A t_2)$  is 1. If  $t_1 \leq_A t_2$  is predicted by  $\mathcal{M}_{\text{rank}}$ , its confidence is from 0 to 1. We only consider  $(t_1, t_2)$  predicted by  $\mathcal{M}_{\text{rank}}$  with  $\text{conf}(t_1 \leq_A t_2) \geq \delta$ , where  $\delta$  is a predefined threshold, as *candidate pairs* to be extended to  $\leq_A$ . We denote the set of all candidate pairs of  $\leq_A$  by  $\leq_A^M$ .

The input and output of the creator are as follows:

- *Input*: A temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  and the augmented training data  $D_{\text{aug}}$  from the critic.
- *Output*: An extended temporal instance  $D'_t = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T)$  based on  $D_{\text{aug}}$  and the predicted orders  $(\leq^M_{A_1}, \dots, \leq^M_{A_n})$ .

(2) *Critic (Sections 5)*. In this phase, the critic of GATE justifies and deduces more temporally ranked pairs, by applying CCs in  $\Sigma$  via the *chase* [62]. Denote the result of chasing by  $\leq^{\Sigma}_A$ . Depending on the validity of  $\leq^{\Sigma}_A$ , we construct an augmented training data, denoted by  $D_{\text{aug}}$ , containing the ranked pairs justified (resp. conflicts caught by CCs);  $D_{\text{aug}}$  will be fed back to the creator, for the next round of model learning, so that the creator can learn from more unseen data and get higher accuracy iteratively. Specifically,  $D_{\text{aug}}$  is a temporal instance extended with a validity flag  $f_{\text{valid}}$ , i.e.,  $D_{\text{aug}} = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{\text{valid}})$ : (a) if  $f_{\text{valid}}$  is true,  $\leq^{\Sigma}_A$  is valid and all temporal order deduced by the chase will be added to  $D_{\text{aug}}$ ; and (b) otherwise, the chasing is invalid, i.e., both  $(t_1, t_2) \in \leq_A$  and  $(t_2, t_1) \in \leq_A$  are deduced, and either  $t_1 <_A t_2$  and  $t_2 <_A t_1$ . In this case, these two *conflicting* orders will be added to  $D_{\text{aug}}$ , and the creator will be asked to resolve this conflict, by revising its model accordingly.

Formally, the input and output of the critic are as follows:

- *Input*: A temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$ , the predicted temporal orders  $(\leq^M_{A_1}, \dots, \leq^M_{A_n})$  and a set  $\Sigma$  of CCs.
- *Output*: Augmented data  $D_{\text{aug}} = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{\text{valid}})$ .

The novelty of the critic consists of (a) the deduction using the chase, and (b) an efficient algorithm implementing the chase.

**Workflow.** As shown in Figure 3, GATE takes a temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  and a set  $\Sigma$  of CCs discovered offline as input, and it outputs an extended temporal instance  $D'_t = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T)$  with a total order  $\leq'_A$  defined for each attribute  $A$ .

GATE first initializes the augmented training data  $D_{\text{aug}} = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{\text{valid}} = \text{true})$  (Line 1, details omitted) for the first

*Input*: A temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$ , and a set  $\Sigma$  of CCs.  
*Output*: An extended temporal instance  $D'_t = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T)$  such that for all  $i \in [1, n]$ , (a)  $\leq'_{A_i}$  extends  $\leq_{A_i}$  and (b)  $\leq_{A_i}$  is a total order.

1.  $D_{\text{aug}} := \text{Initialize}(D_t, f_{\text{valid}} = \text{true}, \Sigma)$ ; Initialize the ranking model  $\mathcal{M}_{\text{rank}}$ ;
2. **while** true **do**
3.     /\* The Creator of GATE \*/
4.      $D_t := \text{Extend}(D_t, D_{\text{aug}})$ ;
5.     Train  $\mathcal{M}_{\text{rank}}$  incrementally based on  $D_t$ ;
6.      $(\leq^M_{A_1}, \dots, \leq^M_{A_n}) :=$  the predicted temporal orders by  $\mathcal{M}_{\text{rank}}$ ;
7.     /\* The Critic of GATE \*/
8.      $(\leq^{\Sigma}_{A_1}, \dots, \leq^{\Sigma}_{A_n}) := \text{Chase}(D_t, \Sigma)$ ;
9.      $D_{\text{aug}} := \text{ConstructAugmented}(D, \leq^M_{A_1}, \dots, \leq^M_{A_n}, \leq^{\Sigma}_{A_1}, \dots, \leq^{\Sigma}_{A_n}, T)$ ;
10.    **if**  $D_t$  no longer changes **then**
11.      **break**;
12.     $D_t = \text{Extend}(D_t, \mathcal{M}_{\text{rank}})$ ;
13. **return**  $D'_t$ ;

Figure 3: Workflow of GATE

round of GATE, by deducing its initial temporal orders  $\leq'_A$  via CCs in  $\Sigma$  whose preconditions do not involve timeliness comparison, e.g., we can create temporal orders for those tuples with available timestamps using  $\varphi_3$ . The initialization can be done efficiently by [33].

Then GATE iteratively executes the creator and the critic in *rounds*. In the  $i$ -th round, (a) the creator extends the temporal instance  $D_t$  (Line 4, see Section 4) based on the augmented training data  $D_{\text{aug}}$  returned by the critic in the  $(i-1)$ -th round, by possibly revising its model to resolve the conflicts. Then the creator *incrementally* trains  $\mathcal{M}_{\text{rank}}$  (Line 5, see Section 4) and predicts new temporal orders  $(\leq^M_{A_1}, \dots, \leq^M_{A_n})$  (Line 6), which are candidate pairs  $(t_1, t_2)$  with confidences at least  $\delta$ ; and (b) the critic deduces more ranked pairs  $(\leq^{\Sigma}_{A_1}, \dots, \leq^{\Sigma}_{A_n})$  by chasing with CCs in  $\Sigma$  (Line 8). Based on the result of chasing, the critic constructs augmented training data  $D_{\text{aug}}$  via procedure *ConstructAugmented* (Line 9, see Section 5); this  $D_{\text{aug}}$  is fed back to the creator for the next round.

Finally, when  $D_t$  no longer changes (Line 10-11), the iteration ends. If  $D_t$  is still not a temporal instance with total orders defined on all attributes, we extend it using procedure *Extend* (Line 12), such that for each pair  $(t_1, t_2)$ , one of  $(t_1, t_2)$  and  $(t_2, t_1)$  learned with a higher confidence is in  $\leq_A$ , until each  $\leq_A$  becomes a total order.

**Termination.** We prove in [1] that eventually, GATE will terminate (Line 10), i.e., with more iterations,  $D_t$  is gradually extended with more orders and finally, becomes stable and does not change.

**Theorem 1:** GATE is guaranteed to terminate.  $\square$

**Example 2:** Continuing with Example 1, assume that  $\Sigma = \{\varphi_1\text{-}\varphi_6\}$  and  $D_t$  has empty temporal orders. We first initialize the training data  $D_{\text{aug}}$  by applying CCs in  $\Sigma$  that do not compare timeliness in their preconditions, e.g., we can deduce  $t_5 \leq_{\text{job}} t_6$  by  $\varphi_3$ . Suppose that after training  $\mathcal{M}_{\text{rank}}$  based on the initial  $D_{\text{aug}}$ , no tuple pair predicted by  $\mathcal{M}_{\text{rank}}$  is at least  $\delta$ -confident in the first round. The creator outputs empty  $\leq^M_A$  for each  $A$  due to the lack of information. Then by applying the CCs in  $\Sigma$ , the critic can deduce new temporal orders  $\leq^{\Sigma}_A$ , e.g.,  $t_1 \leq_{\text{address}} t_2$  by  $\varphi_4$  and  $t_1 \leq_{\text{status}} t_4$  by  $\varphi_5$ . Both  $\leq^M_A$  and  $\leq^{\Sigma}_A$  will be used to construct the augmented training data  $D_{\text{aug}}$ , based on which the creator extends  $D_t$  and incrementally trains  $\mathcal{M}_{\text{rank}}$  in the second round. This time the creator might be able to predict confident temporal orders, e.g.,  $t_1 \leq_{\text{LN}} t_2$ , based on the augmented training data  $D_{\text{aug}}$ . The iteration continues until

$D_t$  does not change anymore. Each temporal order  $\leq_A$  in  $D_t$  will be extended to a total order by  $\mathcal{M}_{\text{rank}}$  if it is still not total.  $\square$

*Remark.* We adopt a confidence threshold  $\delta$  to ensure the reliability of ML predictions, so that only reliable orders are considered. When confident orders cannot be decided for the lack of initial information, we may opt to invite user inspection to ensure the correctness of a few initial ranked pairs, from which more reliable orders can be iteratively deduced/learned. The parameter  $\delta$  plays an important role in model  $\mathcal{M}_{\text{rank}}$ , e.g., when testing  $\mathcal{M}_{\text{rank}}$  on procedural billing data (with 82.28% *real* missing timestamps [40]), we find that 22.6% pairs are identified as confident when  $\delta = 0.55$ . The percentage decreases to 16.7% when  $\delta = 0.6$ . We will test the impact of  $\delta$  in Section 6.

## 4 CREATOR

In this section, we develop the creator of GATE. Given a temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  and the augmented training data  $D_{\text{aug}}$  (from the critic), the creator (a) extends each  $\leq_A$  of  $D_t$  based on  $D_{\text{aug}}$  and (b) trains a ranking model  $\mathcal{M}_{\text{rank}}$  to predict new orders.

**Review on ranking models.** *Learning to Rank* [50] aims to learn a ranking model so that objects can be ranked based on their degrees of relevance, preference, or importance (in our setting, timeliness).

We adopt pairwise ranking since (a) its semantics is consistent with temporal orders, which is a set of *tuple pairs* on which partial orders are defined (Section 2); and (b) by transitivity of temporal orders, pairwise ranking helps us get a total order for all attributes.

**Challenges.** One naturally wants to adopt an existing model. However, it is hard to directly apply one, since the unique features of temporal orders are not well considered, resulting in poor performance.

*(1) Attribute correlation.* Due to the correlated nature of temporal order, we often need to reference other attributes to determine the orders of a given attribute. Moreover, since a value may change back and forth (e.g., the marital status changes from “married” to “divorced”, and back from “divorced” to “married”), it is hard to determine the up-to-date value based on a single attribute only.

*(2) Limitation of embedding models.* To determine the timeliness, care should be taken for lexically different but semantically similar values (e.g., “baby” vs. “birth” and “dead” vs. “expired” for attribute status). Although existing embedding models (e.g., ELMo [59] or Bert [15]) are widely adopted, they cannot be directly used here since they are not trained to organize data chronologically.

*(3) Adaptive margin.* Existing ranking strategies do not consider real-life characteristics of timeliness, e.g., the timespan for a person’s status to move from “birth” to “engaged” is typically longer than from “engaged” to “married” (Figure 4). Instead of ranking status with a *fixed* margin as most existing strategies did, we need a new methodology to embed values using *adaptive* margins, to conform to their real-life behaviors and justify the semantic of ranking.

**Model overview.** We propose a ranking model to tackle the above challenges, whose novelty includes (a) a context-aware scheme that embeds each value along with other correlated values, (b) an encoding mechanism to re-organize the embeddings in a chronological manner, and (c) an attribute-centric adaptive ranking strategy.

As shown in Figure 4, our ranking model  $\mathcal{M}_{\text{rank}}$  takes the current  $D_t$  as input, and outputs new ranked pairs, where each  $(t_1, t_2)$

is associated with a *confidence*, indicating how likely  $t_1 \leq_A t_2$  holds.

Our ranking model consists of three stages as follows: context-aware embedding, chronological encoding and order prediction.

(1)  $\mathcal{M}_{\text{rank}}$  first builds a context-aware embedding for each attribute value using pre-trained language models (ELMo [59] or Bert [15]), in a way that information of correlated values is also embedded.

(2) Based on the embeddings,  $\mathcal{M}_{\text{rank}}$  encodes a target vector  $\phi_A$  for attribute  $A$ , via *non-linear transformation* (see below). Similarly, a value vector  $\phi_{t[A]}$  is encoded for each  $A$ -attribute value of tuple  $t$ , so that (a) if  $t[A]$  is more current,  $\phi_{t[A]}$  is closer to  $\phi_A$  and (b) the gap between  $\phi_{t[A]}$  and  $\phi_A$  is trained adaptively, to reflect real semantics.

(3) Finally, given  $\phi_{t_1[A]}$  and  $\phi_{t_2[A]}$  of  $t_1$  and  $t_2$ ,  $\mathcal{M}_{\text{rank}}$  predicts the order, i.e., whether  $t_1 \leq_A t_2$  holds with high enough confidence.

We next briefly elaborate the context-aware embedding and the chronological encoding scheme with the adaptive margin.

*Context-aware embedding.* To reference correlated values, we treat tuples as sequences and adopt the idea of serialization [47] (so that tuples can be meaningfully ingested by models) to embed values.

Following [47], given a tuple  $t$  in  $D_t$ , we serialize its values:

$$\text{serialize}(t) = \langle \text{COL} \rangle_{A_1} \langle \text{VAL} \rangle_{t[A_1]} \dots \langle \text{COL} \rangle_{A_n} \langle \text{VAL} \rangle_{t[A_n]},$$

where  $\langle \text{COL} \rangle$  and  $\langle \text{VAL} \rangle$  are special tokens, denoting the start of attribute and value, respectively (see Figure 4). This serialization is fed as input to a pre-trained language model  $\text{emb}(\cdot)$  to compute a  $d$ -dimensional embedding for each  $A$ -attribute value, denoted by  $\text{emb}(t[A]) \in \mathbf{R}^d$ . Besides, we average out the embedding vectors for all  $t[A]$  to get a context representation of tuple  $t$ , i.e.,  $\text{emb}(t) = \frac{1}{n} \sum_{i=1}^n \text{emb}(t[A_i])$ . Finally for each  $A$ -attribute value of  $t$ , we get the context-aware embedding of  $t[A]$ , denoted by  $E_{t[A]} \in \mathbf{R}^{2d}$ :

$$E_{t[A]} = [\text{emb}(t[A]); \text{emb}(t)],$$

where  $[\cdot]$  denotes vector concatenation. In this way,  $E_{t[A]}$  embeds not only the  $A$ -attribute value, but also the contextual information from other attribute values, to allow comprehensive ranking. Similarly, a schema embedding for each attribute  $A$  is computed:  $E_A = \text{emb}(A)$ , by feeding the attribute name, e.g., status, to the model.

*Chronological encoding with adaptive margins.* While pre-trained embeddings are widely adopted to capture semantics, they are not trained to organize temporal orders. Thus we propose chronological encoding to re-organize the embeddings to preserve timeliness. The idea is to use schema embedding as the target and make the embedding of a more current value closer to the target; moreover, instead of ranking in fixed margins, embeddings are ordered adaptively.

Specifically, given the embedding of the  $A$ -attribute of  $t$ , i.e.,  $E_{t[A]}$ , we encode it using a context encoder  $\text{ENC}_{\text{ctx}}(\cdot)$  as follows:

$$\phi_{t[A]} = \text{ENC}_{\text{ctx}}(E_{t[A]}) = \sigma(W_2 * \sigma(W_1 * E_{t[A]})),$$

where  $W_1$  and  $W_2$  are learnable parameters of the encoder, and  $\sigma$  is the non-linear sigmoid activation function given by  $\sigma(x) = \frac{1}{1+e^{-x}}$ .

Similarly, the target vector for attribute  $A$  is encoded as  $\phi_A = \text{ENC}_{\text{attr}}(E_A)$ , where  $\text{ENC}_{\text{attr}}(\cdot)$  denotes the schema encoder.

To train the encoders with ordered embeddings and adaptive margins, we adopt an attribute-centric adaptive margin-based triplet loss. Given  $\leq_A$  in  $D_t$ , the loss on  $A$  is formulated as follows:

$$\text{loss}(A) = \sum_{(t_1, t_2) \in \leq_A} \left\{ \max\{-\tanh(\langle \phi_{t_2[A]}, \phi_A \rangle) \right\}$$

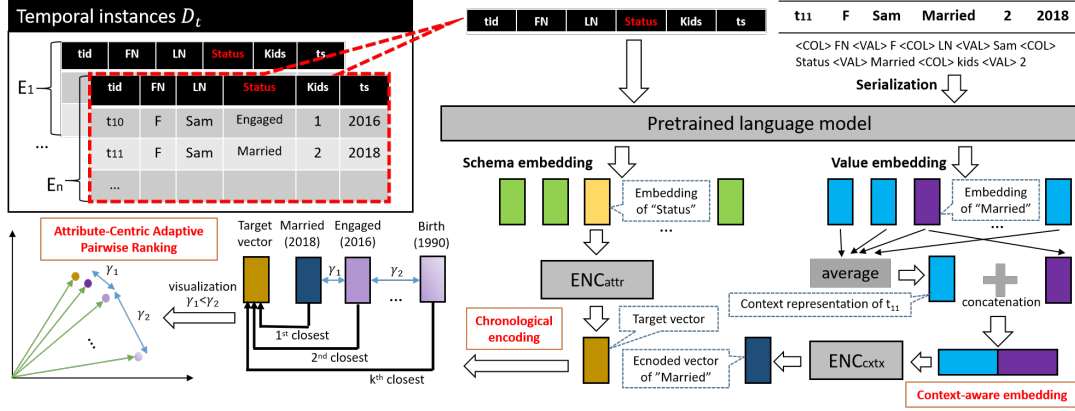


Figure 4: The Network Architecture of Creator

$$+\gamma_{t_1, t_2} + \tanh(\langle \phi_{t_1[A]}, \phi_A \rangle, 0),$$

where  $\langle \cdot, \cdot \rangle$  is the inner product and  $\gamma_{t_1, t_2}$  is the adaptive margin between the two tuples; we set  $\gamma_{t_1, t_2}$  to be  $1 + \cos(\phi_{t_1[A]}, \phi_{t_2[A]})$  in practice, which characterize the co-occurrence of  $t_1[A]$  and  $t_2[A]$ .

Intuitively, by minimizing the loss, for each training instance  $t_1 \leq_A t_2$ , (a) we make the encoded vector of the more current value  $t_2[A]$  closer to target  $\phi_A$  (the first term) and (b) we ensure an adaptive margin  $\gamma_{t_1, t_2}$  between the two tuples (the second term). In other words, the attribute values in the encoded space are not only arranged chronologically by their “distances” to  $\phi_A$ , from which temporal orders can be easily derived, their margins are also adaptively determined, to reflect the semantic of timeliness ranking.

**Model training and instance extension.** In each round, the creator receives augmented training data  $D_{aug} = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{valid})$ , based on which it (a) extends the temporal instance  $D_t$  with  $D_{aug}$  and (b) incrementally trains  $\mathcal{M}_{rank}$  via back-propagation.

In the first round,  $D_{aug}$  is initialized to be the temporal orders constructed by applying those CCs in  $\Sigma$  without timeliness comparison in their preconditions [33]. In the following rounds,  $D_{aug}$  is constructed by the critic, based on the result of chasing with CCs.

Specifically, depending on flag  $f_{valid}$  in  $D_{aug}$ , we have two cases:

(1) If  $f_{valid}$  is true, the result of chasing is valid and  $\mathcal{M}_{rank}$  is incrementally trained based on the orders in  $D_{aug}$  (see below). Moreover, the temporal instance  $D_t = (D, \leq_{A_1}, \dots, \leq_{A_n}, T)$  is extended with the ranked pairs in  $D_{aug}$ , *i.e.*, for each  $(t_1, t_2)$  in  $\leq'_A$  of  $D_{aug}$ ,  $(t_1, t_2)$  is added to  $\leq_A$ . In this case, we say that  $(t_1, t_2)$  becomes *stable*. Once a temporal order becomes stable, it will not be removed from  $D_t$ .

(2) If  $f_{valid}$  is false, the result of chasing is invalid, *i.e.*, there is an attribute  $A$  such that conflicting orders  $(t_1, t_2)$  and  $(t_2, t_1)$  are both in  $\leq'_A$  of  $D_{aug}$  *i.e.*,  $\{(t_1, t_2), (t_2, t_1)\} \subseteq \leq'_A$ , and either  $t_1 <_A t_2$  or  $t_2 <_A t_1$ . In this case, we decide that either  $(t_1, t_2)$  or  $(t_2, t_1)$  is added to  $\leq_A$  using  $\mathcal{M}_{rank}$ , with a higher confidence (and possibly user inspection). Assume *w.l.o.g.* that  $(t_1, t_2)$  is added to  $\leq_A$  (*i.e.*, it becomes stable). Then, the creator fine-tunes  $\mathcal{M}_{rank}$  so that  $t_1$  and  $t_2$  are better separated in the encoded space (see [1] for details).

**Incremental training.** Incremental training of  $\mathcal{M}_{rank}$  might lead to the catastrophic forgetting issue [39], *i.e.*, the model might forget some temporal orders learned in prior rounds. To overcome this, we adopt a simple strategy to retain prior knowledge: In each round,  $D_t$  is augmented with  $D_{aug}$  and  $\mathcal{M}_{rank}$  is incrementally trained on

$D_t$  from the last checkpoint. In this way, the model is able to learn from previous training instances and avoid catastrophic forgetting.

**Monotonicity.** One can verify that the number of stable temporal orders in  $D_t$  is (strictly) monotonically increasing when more rounds GATE are performed (see a detailed lemma in [1]). The termination of GATE (Theorem 1) partly depends on this monotonicity.

**Confidence.** Given  $\leq_A$ , a tuple pair  $(t_1, t_2)$  and a closeness function  $f_L$ ,  $\mathcal{M}_{rank}$  predicts  $(t_1, t_2) \in \leq_A$  if  $f_L(\phi_{t_2[A]}, \phi_A) > f_L(\phi_{t_1[A]}, \phi_A)$ ; its confidence indicates how likely  $t_1 \leq_A t_2$  holds, *i.e.*, we have

$$\text{conf}(t_1 \leq_A t_2) = \sigma(f_L(\phi_{t_2[A]}, \phi_A) - f_L(\phi_{t_1[A]}, \phi_A))$$

and set  $\text{conf}(t_2 \leq_A t_1)$  to 0. Thus, given a confidence threshold  $\delta > 0$ , we will not predict both  $t_1 \leq_A t_2$  and  $t_2 \leq_A t_1$  as confident orders.

We use the negative Euclidean distance  $-L_2(\phi_{t_1[A]}, \phi_A)$  in practice to measure the “closeness” between  $\phi_{t_1[A]}$  and  $\phi_A$ , where the one with a larger closeness value is more current. The distance gap between  $\phi_{t_1[A]}$  and  $\phi_{t_2[A]}$  quantifies the confidence: the larger the gap, the larger the confidence, which ranges from 0 to 1.

**Example 3:** Consider the example in Figure 4, where we focus on attribute status. After creating the context-aware embedding based on a pre-trained model, it chronologically encodes a target vector  $\phi_{status}$  and value vectors for all values, so that they are arranged by their distances to  $\phi_{status}$ . To illustrate, we also label the *unknown* timestamp of each value vector in the figure (*e.g.*,  $T_e(\text{Married}) = 2018$ ). Since  $\phi_{\text{Married}}$  is the closest to  $\phi_{status}$ , it is predicted to be the latest status value and new ranked pairs are constructed accordingly for  $\leq_{status}$ , as augmented training data in the next round.  $\square$

**Remark.** Our creator learns temporal orders by utilizing context-aware embedding, chronological encoding and attribute-centric adaptive ranking. However, it does not explicitly take into account of some temporal properties, such as the transitivity. This motivates us to use critic to deduce and justify the orders in the next section.

## 5 CRITIC

In this section, we develop the critic under GATE for justifying and deducing temporal orders. Taking a temporal instance  $D_t$ , the temporal orders  $(\leq_{A_1}^M, \dots, \leq_{A_n}^M)$  predicted by the creator and a set  $\Sigma$  of mined CCs as input, the critic (a) deduces more ranked pairs  $(\leq_{A_1}^\Sigma, \dots, \leq_{A_n}^\Sigma)$  by applying CCs via the *chase*, and (b) constructs the augmented training data  $D_{aug}$  and feeds  $D_{aug}$  back to the creator.

## 5.1 Chasing with CCs

We extend the classic chase in [32] for CCs under GATE. In the following, we first specify fixes and ground truth. We then present the chase for CCs, with the Church-Rosser property.

**Fixes.** We extend temporal orders in  $D_t$  by applying CCs in  $\Sigma$  to deduce *fixes*, which are modeled as sets of ranked pairs, denoted by  $\bar{U} = (\bar{U}_{A_1}, \dots, \bar{U}_{A_n})$ , where each ranked pair  $(t_1, t_2)$  in  $\bar{U}_A$  is referred to as a *fix* and it means that  $t_1 \leq_A t_2$  or  $t_1 <_A t_2$  is deduced. We only apply a CC if its precondition is satisfied by a collection  $\Gamma$  of “ground truth”. Intuitively, the fixes are logical consequences of  $\Sigma$  and  $\Gamma$ , *i.e.*, as long as the CCs in  $\Sigma$  and  $\Gamma$  are correct, so are the fixes.

**Validity.** We say  $\bar{U}$  is *valid* if it has no conflicting fixes, *i.e.*, there exist no attribute  $A$  and tuples  $t_1, t_2$  such that  $(t_1, t_2) \in \bar{U}_A$  and  $(t_2, t_1) \in \bar{U}_A$  at the same time, and either  $t_1 <_A t_2$  or  $t_2 <_A t_1$ .

**Ground truth.** To justify the correctness of fixes, we maintain and employ a collection  $\Gamma = (\Gamma_{A_1}, \dots, \Gamma_{A_n})$  of validated data, *enclosed* in  $\bar{U}$ . In our setting, block  $\Gamma$  is initialized by applying CCs in  $\Sigma$  whose preconditions do not involve timeliness comparison (*e.g.*, initial temporal orders with partial timestamps via  $\varphi_3$ ), and is iteratively expanded with the temporal orders learned by the creator with confidence above threshold  $\delta$  or deduced by the chase in the critic.

**The chase.** Given a temporal instance  $D_t$ , the chase deduces fixes by chasing  $D_t$  with CCs in  $\Sigma$  and ground truth in  $\Gamma$ . It uses sets  $\leq^{\Sigma} = (\leq^{\Sigma}_{A_1}, \dots, \leq^{\Sigma}_{A_n})$  to keep track of the affected fixes in  $\bar{U}$ . Specifically, the  $i$ -th chase step of  $D_t$  by  $\Sigma$  at  $(\bar{U}_i, \leq^{\Sigma}_i)$  is:

$$(\bar{U}_i, \leq^{\Sigma}_i) \Rightarrow_{(\varphi, h)} (\bar{U}_{i+1}, \leq^{\Sigma}_{i+1}),$$

where  $\varphi : X \rightarrow p_0$  is a CC in  $\Sigma$ ,  $h$  is a valuation of  $\varphi$  in  $\mathcal{D}_t$ , and the application of  $(\varphi, h)$  should satisfy the following conditions:

- (1) All predicates  $p \in X$  are *validated*, *i.e.*, if  $p$  is  $t[A] \oplus c$  or  $t_1[A] \oplus t_2[A]$ , then  $h \models p$ ; and if  $p$  is  $t_1 \leq_A t_2$ , then  $(t_1, t_2)$  is in  $\bar{U}_A$ .
- (2) The consequence  $p_0 : t_1 \leq_A t_2$  extends  $\bar{U}_i$  to  $\bar{U}_{i+1}$ , such that  $(t_1, t_2)$  is added to  $\bar{U}_A$  of  $\bar{U}_i$ ; similarly,  $p_0$  extends  $\leq^{\Sigma}_i$  to  $\leq^{\Sigma}_{i+1}$ .

**Chasing.** Starting from a set  $\bar{U}_0$  of fixes, initialized to be  $\Gamma$ , and an empty  $\leq^{\Sigma}_0$ , a chasing sequence  $\xi$  of  $D_t$  by  $(\Sigma, \Gamma)$  is

$$(\bar{U}_0, \leq^{\Sigma}_0), \dots, (\bar{U}_k, \leq^{\Sigma}_k),$$

where  $(\bar{U}_i, \leq^{\Sigma}_i) \Rightarrow_{(\varphi, h)} (\bar{U}_{i+1}, \leq^{\Sigma}_{i+1})$  is a valid chase step, *i.e.*, a valuation  $h$  of  $\varphi$  extends  $(\bar{U}_i, \leq^{\Sigma}_i)$  to  $(\bar{U}_{i+1}, \leq^{\Sigma}_{i+1})$  where  $\bar{U}_{i+1}$  is valid.

The chasing sequence is *terminal* if there exist no CC  $\varphi$  in  $\Sigma$  and valuation  $h$  of  $\varphi$  such that  $(\varphi, h)$  leads to another valid chase step.

A chase sequence  $\xi$  terminates in one of the following cases:

- (1) No more CCs in  $\Sigma$  can be applied. If so, we say that  $\xi$  is valid, with  $(\bar{U}_k, \leq^{\Sigma}_k)$  as its result.
- (2) Either  $\bar{U}_0$  is invalid or there exist  $\varphi, h, \bar{U}_{k+1}$  and  $\leq^{\Sigma}_{k+1}$  such that  $(\bar{U}_k, \leq^{\Sigma}_k) \Rightarrow_{(\varphi, h)} (\bar{U}_{k+1}, \leq^{\Sigma}_{k+1})$  but  $\bar{U}_{k+1}$  is invalid. Such  $\xi$  is invalid, and the result of the chase is  $\perp$  (undefined).

Intuitively, the chase helps us deduce more ranked pairs when it terminates with enriched  $(\bar{U}_k, \leq^{\Sigma}_k)$ ; moreover, it justifies and explains the learned order if no invalid chase step is taken. When its result is  $\perp$ , it detects invalid ranked pairs of the learner.

**Example 4:** Consider  $D_t$  in Figure 1. Assume that  $\Sigma = \{\varphi_1 - \varphi_6\}$  and  $\leq^{\Sigma} = (\leq^{\Sigma}_{A_1}, \dots, \leq^{\Sigma}_{A_n})$ , where each  $\leq^{\Sigma}_{A_i}$  is empty. We initialize  $\bar{U}_0$  and  $\Gamma$  by applying CCs without timeliness comparison, as we did in Example 2, *e.g.*, since  $t_1$  [status] (resp.  $t_2$  [status]) is “single” (resp. “married”) in Figure 1,  $t_1 \leq_{\text{status}} t_2$  is initialized in  $\Gamma$  by applying  $\varphi_2$ .

We have the following chase steps of  $D_t$  by  $(\Sigma, \Gamma)$ :

- (1) By applying  $(\varphi_4, h_4)$ , where  $\varphi_4$  is  $t_1 \leq_{\text{status}} t_2 \rightarrow t_1 \leq_{\text{address}} t_2$  and  $h_4$  maps the variables of  $\varphi_4$  to tuples  $t_1$  and  $t_2$  in  $D_t$ , we deduce  $t_1 \leq_{\text{address}} t_2$  by the chase step  $(\bar{U}_0, \leq^{\Sigma}_0) \Rightarrow_{(\varphi_4, h_4)} (\bar{U}_1, \leq^{\Sigma}_1)$ , *i.e.*,  $\bar{U}_1$  extends  $\bar{U}_0$  by adding  $(t_1, t_2)$  to  $\bar{U}_{\text{address}}$ ; similar to  $\leq^{\Sigma}_1$ .
- (2) The chase proceeds to deduce  $t_1 \leq_{\text{status}} t_4$  by applying  $\varphi_5$ .

This chasing sequence is valid since each chase step in the sequence is valid and no more CCs in  $\Sigma$  can be applied anymore.  $\square$

**Church-Rosser property.** Following [2], we say that chasing with CCs is *Church-Rosser* if for any temporal instance  $D_t$ , any set  $\Sigma$  of CCs, any collection  $\Gamma$  of ground truth, all chasing sequences of  $D_t$  by  $(\Sigma, \Gamma)$  are terminal and converge at the same result. Below we show that chasing with CCs is Church-Rosser (proven in [1]).

**Corollary 2:** *Chasing with CCs is Church-Rosser.*  $\square$

## 5.2 Deduction with the Chase

No matter how desirable, the chase could be expensive if we enumerate valuations of CCs in an exhaustive manner. Below we provide an efficient algorithm to implement the chase.

**Challenges.** A brute-force implementation of the chase is by enumerating valuations  $h$  of each CC  $\varphi$  in  $\Sigma$ . If  $(\varphi, h)$  can be applied, a chase step is performed, until the chasing sequence terminates. This method is, however, costly since valuation enumeration is inherently exponential. To tackle this challenge, we develop an efficient algorithm to implement the chase; the key idea is to only evoke valuations *pertaining* to the affected fixes in the chase *lazily* (see below).

We assume *w.l.o.g.* that for each  $\varphi : X \rightarrow p_0$  in  $\Sigma$ ,  $X$  has a predicate  $p$  in the form of  $t_1 \leq_A t_2$  (*e.g.*,  $\varphi_4$ ). For those CCs that do not compare timeliness in their precondition (*e.g.*,  $\varphi_1$ ), we apply them in a pre-processing step, to generate initial temporal orders in  $D_t$ .

**Lazy evocation.** To allow lazy evocation, the valuations of CCs in  $\Sigma$  are generated only when they are evoked by some newly deduced orders, instead of constructing all at the beginning of the chase.

Specifically, when a new temporal order  $o$  is deduced, we check each  $\varphi : X \rightarrow p_0$  in  $\Sigma$  and evoke a new valuation  $h$  of  $\varphi$  if (a)  $o$  corresponds to a predicate in  $X$  (*i.e.*,  $o$  is validated in  $h$ ); in this case, we say that  $h$  is a valuation *pertaining* to  $o$  since  $h$  is “activated” by  $o$ , (b)  $h$  has not been evoked before and (c) the order that  $h$  deduces, *i.e.*,  $h(p_0)$ , is not deduced by other valuations before. We maintain designated data structures for checking conditions (a), (b) and (c) efficiently. See [1] for a detailed description of the data structures.

**Algorithm.** Putting these together, we present Chase in Figure 5. It returns new orders  $\leq^{\Sigma}$  if the chase is valid, and  $\perp$  otherwise.

Chase starts with the initialization (Line 1). (a) Ground truth  $\Gamma$  is initialized with all stable ranked pair in  $D_t$  via procedure Initialize (omitted). (b) It initializes  $\bar{U}$  and  $\leq^{\Sigma}$  as stated in Section 5.1 to be  $\Gamma$  and  $\emptyset$ , respectively. (c) The set  $\Delta$  of newly validated orders is initialized to be the newly stable orders in  $\Gamma$  via procedure NewStable



(omitted); and they are the temporal orders “triggering” the chase.  
(d) The set  $\mathcal{H}$  of evoked valuations is initialized to be empty.

Then for each order  $o$  in  $\Delta$ , Chase does the following (Line 5-15):  
(a) Evoke the valuations pertaining to  $o$  via the lazy evocation strategy stated above, by calling CCEvoke (omitted), and add them to  $\mathcal{H}$  (Line 5). (b) For each valuation  $h \in \mathcal{H}$  pertaining to  $o$  that deduces unknown ranked pair (checked in Line 7-8), mark  $o$  as validated in  $h$  (Line 9). If all predicates in the precondition of  $h$  are validated (Line 10),  $(h, \varphi)$  can be applied and the consequence  $t_1 \leq_A t_2$  of  $h$  is deduced (Line 11). The ranked pair  $(t_1, t_2)$  is added to  $\bar{U}_A$  and  $\leq_A^\Sigma$  (Line 12). (c) Check conflicts (Line 13-14): if  $(t_1, t_2)$  conflicts with  $(t_2, t_1)$  that is already in  $\bar{U}_A$  or  $\leq_A^M$ , the chase terminates with  $\leq^\Sigma = \perp$ . In this case, the valuations in  $\mathcal{H}$  are kept temporally so that they can be re-used in the next round of GATE when the conflicts are resolved (not shown). (d) Add the newly deduced order to the set  $\Delta^{\text{new}}$  (Line 15) for iterative processing, by assigning  $\Delta^{\text{new}}$  to  $\Delta$  (Line 16).

Finally, the result of chasing,  $\leq^\Sigma$ , is returned (Line 17).

**Example 5:** Recall that  $\varphi_4$  is  $t_1 \leq_{\text{status}} t_2 \rightarrow t_1 \leq_{\text{address}} t_2$  and  $\varphi_5$  is  $t_1 \leq_{\text{status}} t_2 \wedge t_2 \leq_{\text{status}} t_3 \rightarrow t_1 \leq_{\text{status}} t_3$ . Let  $\Sigma = \{\varphi_4, \varphi_5\}$  and  $\Delta = \{t_1 \leq_{\text{status}} t_2\}$ . We process  $t_1 \leq_{\text{status}} t_2$  in  $\Delta$  as follows. It first evokes valuations  $h_4$  and  $h_5$  which map the variables of  $\varphi_4$  and  $\varphi_5$  to tuples  $t_1$  and  $t_2$  in  $D_t$ , respectively, with  $t_1 \leq_{\text{status}} t_2$  validated. Since the predicate  $t_2 \leq_{\text{status}} t_3$  in  $h_5$  is not validated,  $h_5$  is kept in  $\mathcal{H}$  for later processing. In contrast, all predicates in the precondition of  $h_4$  are validated and  $\varphi_4$  deduces  $t_1 \leq_{\text{address}} t_2$ . Suppose that there is no conflicting order in  $\bar{U}_{\text{address}}$  and  $\leq_{\text{address}}^M$ . Then  $t_1 \leq_{\text{address}} t_2$  forms a new set  $\Delta$  and the process continues, until  $\Delta$  is empty.  $\square$

**Complexity.** The loop of Chase executes at most  $O(|R||D|^2)$  times, since there are at most  $|R||D|^2$  temporal orders to be deduced. For each temporal order  $o$  deduced, we evoke CCs based on  $o$  and update data structures in  $O(c_{\text{val}}|\Sigma|)$  time, where  $c_{\text{val}}$  denotes the unit cost of constructing valuations for fixed  $o$  and  $\varphi$ , and there are at most  $|\Sigma|$  CCs. Thus Chase takes at most  $O(c_{\text{val}}|\Sigma||R||D|^2)$  time.

**Augmented training data construction.** Recall that the result of chasing, denoted by  $\leq^\Sigma$ , is valid or invalid. Based on  $\leq^\Sigma$ , we construct the augmented training data  $D_{\text{aug}}$  as follows.

(1) If  $\leq^\Sigma$  is valid, both the temporal orders deduced by the chase and predicted by the creator are used to create  $D_{\text{aug}} = (D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{\text{valid}} = \text{true})$  where  $\leq'_{A_i} = \leq_{A_i}^M \cup \leq_{A_i}^\Sigma$  ( $i \in [1, n]$ ).

(2) If  $\leq^\Sigma$  is  $\perp$ , then there exist conflicting ranked pairs, *i.e.*, both  $(t_1, t_2)$  and  $(t_2, t_1)$  are in  $\bar{U}_A$  or  $\leq_A^M$  with  $t_1 <_A t_2$  or  $t_2 <_A t_1$ . In this case, we construct  $D_{\text{aug}}$  to be  $(D, \leq'_{A_1}, \dots, \leq'_{A_n}, T, f_{\text{valid}} = \text{false})$  where  $\leq_{A_i} = \{(t_1, t_2), (t_2, t_1)\}$  if  $A_i = A$  and  $\leq_{A_i} = \emptyset$  otherwise.

As shown in Section 4, the creator fine-tunes its model by using the deduced ranked pairs or the detected conflicts in  $D_{\text{aug}}$ .

## 6 EXPERIMENTAL STUDY

Using real-life and synthetic data, we evaluated (1) the effectiveness and (2) the efficiency of GATE for determining temporal orders. We also (3) conducted a case study to showcase the usefulness of GATE.

**Experimental settings.** We start with the experimental setting.

**Datasets.** We used three real-life datasets and one synthetic dataset.

*Input:* A temporal instance  $D_t$ , the set  $\Sigma$  of CCs, the predicted  $(\leq_{A_1}^M, \dots, \leq_{A_n}^M)$   
*Output:* The result of chasing,  $\leq^\Sigma$ .

```

1.  $\Gamma := \text{Initialize}(D_t); \bar{U} := \Gamma; \leq^\Sigma := \emptyset; \Delta = \text{NewStable}(\Gamma); \mathcal{H} := \emptyset;$ 
2. while  $\Delta$  is not empty do
3.    $\Delta^{\text{new}} = \emptyset;$ 
4.   for each  $o \in \Delta$  do
5.      $\mathcal{H} := \mathcal{H} \cup \text{CCEvoke}(D_t, \Sigma, o);$ 
6.     for each  $h$  of  $\varphi : X \rightarrow t_1 \leq_A t_2$  in  $\mathcal{H}$  s.t.  $h$  pertains to  $o$  do
7.       if the order between  $t_1[A]$  and  $t_2[A]$  is already settled then
8.          $\mathcal{H} := \mathcal{H} \setminus \{h\}$ ; continue ;
9.       Mark  $o$  as validated in  $h$ ;
10.      if all predicates in the precondition of  $h$  are validated then
11.         $\mathcal{H} := \mathcal{H} \setminus \{h\}$ ; /*  $t_1 \leq_A t_2$  is a newly deduced order */
12.         $\bar{U}_A := \bar{U}_A \cup (t_1, t_2)$ ;  $\leq_A^\Sigma := \leq_A^\Sigma \cup (t_1, t_2)$ ;
13.        if  $(t_2, t_1) \in \bar{U}_A$  or  $\leq_A^M$  with  $t_1 <_A t_2$  or  $t_2 <_A t_1$  then
14.           $\leq^\Sigma := \perp$ ; return  $\leq^\Sigma$ ;
15.           $\Delta^{\text{new}} := \Delta^{\text{new}} \cup \{t_1 \leq_A t_2\}$ ;
16.         $\Delta := \Delta^{\text{new}}$ ;
17.      return  $\leq^\Sigma$ ;

```

Figure 5: Procedure Chase

(1) Career [41], a benchmark about the careers of football players from FIFA-15 to FIFA-22; it contains 108.5K tuples from 27.2K entities with 20 attributes. We determine the timeliness of potential, position, reputation and league name. (2) NBA [18], a dataset that encompasses the careers of basketball players; it contains 10.6K tuples with 12 attributes. We determine the currency of team, pts (points) and weight. (3) COM [34], an open-source dataset about self-employed entrepreneurs in Shenzhen. After removing duplicated tuples and digital attributes (*e.g.*, the organization code), the dataset has 1,983,698 tuples and 8 attributes. We derive the timeliness of entrepreneur names. (4) Person, a synthetic dataset with 12.3K tuples from 1K entities. Just like Figure 1, we adopted 7 attributes and determine the currency of LN, Status, Kids. Here Person is generated by enforcing CCs (*e.g.*,  $\varphi_1$ - $\varphi_6$ ), to simulate real-world scenarios.

All the datasets have ground truth, *i.e.*, all tuples carry timestamps and are grouped by entities. The timestamps of COM are in seconds, while the others are in years; it is reasonable since, *e.g.*, it is uncommon for NBA players to frequently change teams in one year. We randomly selected 5% data with initial timestamps and masked the remaining. This default ratio *ts%* of initial timestamps is set intentionally small (so that the problem is more challenging); we will test the impact of *ts%* by varying *ts%* (Exp-1).

**Currency constraints.** We extended DCFinder [58] to discover CCs as discussed in Section 2. Note that CC discovery is conducted once on each dataset offline. Besides, we manually checked and adjusted the CCs discovered to ensure their correctness. We found 42, 32, 40 and 36 CCs for Career, NBA, COM and Person, respectively.

**ML model.** To learn the ranking model  $\mathcal{M}_{\text{rank}}$ , we used Bert [15] (distilbert) with 768 dimension to initialize the embeddings. We adopted 2 hidden layers in our encoders with sizes 200 and 100, respectively. The margin  $\gamma$  was adaptively computed and the model was trained with 30 epochs using Adam optimizer [38]. The learning rate is  $1e-4$ . We used 5% data with timestamps as training data.

**Baselines.** GATE was implemented in Python and we compared it with the following baselines: (1) Creator, a variant of GATE with the creator only, *i.e.*, it predicts temporal orders using  $\mathcal{M}_{\text{rank}}$ ; (2)

Critic, a variant of GATE with the critic only, *i.e.*, it deduces temporal orders by chasing with CCs; (3) Creator<sub>itr</sub>, a variant of Creator that iteratively updates its training data with predicted but unjustified temporal orders. (4) Creator<sub>NC</sub>, Creator<sub>NE</sub>, Creator<sub>NA</sub>, another three variants of Creator that implement  $\mathcal{M}_{\text{rank}}$  without contextual information, without chronological encoding, and using regular cross entropy loss instead of adaptive margin-based loss, respectively; (5) GATE<sub>NC</sub>, a variant of GATE that adopts the brute-force method for the chase, by enumerating all valuations exhaustively.

We also tested (6) UncertainRule [44], which uses uncertain currency rules to evaluate data currency; (7) Improve3C [18], a data quality framework that combines completeness, consistency and currency [28]; we only compare its accuracy for currency; (8) RANK<sub>Bert</sub> [56], a state-of-the-art ML ranking model based on Bert; and (9) Ditto<sub>Rank</sub>, a ranking model that first trains a ditto model [47] to conduct binary classification on attribute values (with contextual information) and then sorts all attribute values using ditto as the comparison operator, *i.e.*, we used bubble-sort to sort the values, where the comparison between two values is conducted by ditto.

Among the baselines, (a) Critic, UncertainRule and Improve3C are rule-based, where rules for the latter two are converted from same CCs mined by DCFinder; if total orders cannot be deduced by rules, the unknown temporal orders are randomly decided, (b) RANK<sub>Bert</sub>, Ditto<sub>Rank</sub>, Creator and its variants are ML-based, and (c) GATE<sub>NC</sub> is a hybrid method, which produces same results as GATE. Thus, we compared GATE<sub>NC</sub> mostly for efficiency.

We did the experiments on a single machine powered by 256GB RAM and 32 processors with Intel(R) Xeon(R) Gold 5320 CPU @2.20GHz. We ran each experiment 3 times and report the average.

**Experimental results.** We next report our findings.

**Exp-1: Effectiveness.** Since we adopted the pairwise ranking setting to deduce ranked pairs, we evaluated the accuracy of GATE following [18, 27, 45]: (1) precision, the ratio of temporal orders determined correctly to all ranked pairs predicated true, (2) recall, the ratio of temporal orders predicted correctly to all true orders, and (3)  $F$ -measure  $= 2 \times \frac{\text{recall} \times \text{precision}}{\text{recall} + \text{precision}}$ . To evaluate the ranking of GATE, for each entity instance pertaining to entity  $e$ , we compute  $\text{MRR}(e) = \frac{1}{n} \sum_{i=1}^n \frac{1}{\text{rank}_i}$ , the mean reciprocal rank over a set of  $n$  ranking results, where  $n$  is the number of currency attributes and  $\text{rank}_i$  denotes the rank of the latest value of the  $i$ -th attribute for entity  $e$ , and  $\text{MAP@K}(e) = \frac{1}{n} \sum_{i=1}^n \text{AP@K}(i)$ , the mean average precision at  $K$  that assesses whether the top- $K$  values predicted are relevant and whether the latest  $K$  values are at the top, where  $\text{AP@K}(i)$  is the average precision at  $K$  of the  $i$ -th attribute for entity  $e$ . Let  $D_t$  be a collection of  $k$  entity instances. We report (4)  $\text{MRR} = \frac{1}{k} \sum_{j=1}^k \text{MRR}(e_j)$  and (5)  $\text{MAP@K} = \frac{1}{k} \sum_{j=1}^k \text{MAP@K}(e_j)$ , with  $K = 3$  by default.

**Rounds.** We report the performance of GATE from the first round till its termination; at the end of the fixed round, we use current  $\mathcal{M}_{\text{rank}}$  in the creator. As shown in Figures 6(a)-6(i), GATE takes 11, 12, 7 and 9 rounds to terminate on Career, NBA, COM and Person, respectively, *i.e.*, GATE converges quickly. Besides, we find the following.

(1) Although the performance of GATE might fluctuate (*e.g.*, Figure 6(h)), which is common in ML models [15], all metrics increase with more rounds in most cases, *e.g.*,  $F$ -measure, MRR and

MAP increase from 0.767 to 0.866, 0.786 to 0.857, and 0.752 to 0.809, respectively, after 11 rounds on Career, verifying that GATE is able to deduce the latest values and produce good currency ranking. This is because the creator iteratively accumulates training data from the critic such that the model is better trained with more rounds; meanwhile, with better results predicted by the creator, the critic deduces more orders as augmented training data for the creator in subsequent rounds. Moreover, in Figures 6(b) and 6(c), precision and recall are 0.859 and 0.873, respectively, indicating that GATE achieves a good balance between the two and is fairly accurate. Note that Creator<sub>itr</sub> suffers from the accuracy fluctuation (*e.g.*, Figure 6(i)) since its model is affected by noisy (unjustified) temporal orders accumulated over rounds. The performance of other methods does not depend on rounds, as shown in flat lines. Since GATE behaves similarly under all metrics, below we focus on  $F$ -measure.

(2) On average GATE outperforms Creator and Critic by 7.8% and 34.4% in  $F$ -measure, respectively, up to 11.0% and 45.6%, improving both. The creator and critic benefit each other: (a) Creator produces “hidden” temporal orders for Critic to preform deduction, and (b) Critic deduces and justifies the orders, which are in turn provided as augmented training data to Creator; on average, the critic generates 6K new training data (tuple pairs) per round on COM, improving  $F$ -measure of GATE from 0.701 to 0.748 after 5 rounds.

(3) Creator is more accurate than all its variants, *e.g.*, the average  $F$ -measure of Creator is 0.722, as opposed to 0.641, 0.613 and 0.714 by Creator<sub>NC</sub>, Creator<sub>NE</sub> and Creator<sub>NA</sub>, respectively. Intuitively, (a) without utilizing the contextual information, Creator<sub>NC</sub> cannot reference correlated attributes; (b) Creator<sub>NE</sub> has low accuracy with existing embedding models, and (c) compared to the regular cross entropy loss used in Creator<sub>NA</sub>, the adaptive pairwise ranking loss helps by considering the semantics in ranking. Moreover, GATE is 5.6% more accurate than Creator<sub>itr</sub> on average. This verifies the need for justifying the temporal orders learned by Creator.

(4) The accuracy of GATE is higher than UncertainRule, Improve3C, RANK<sub>Bert</sub> and Ditto<sub>Rank</sub>, *e.g.*, the average  $F$ -measure of GATE is 0.802 as opposed to 0.455, 0.428, 0.659 and 0.651 for the four, respectively. This shows the benefits of combining deep learning and logic rules: (a) compared with rule-based methods, GATE can learn from unseen data and has better generalizability; and (b) compared with ML-based methods, GATE is able to justify the reliability of deduction and produces more training data for the model.

**Varying  $K$ .** We varied parameter  $K$  of MAP@ $K$  from 1 to 5 in Figure 6(f). GATE consistently achieves the highest MAP@ $K$ , *e.g.*, 5% higher than the best baseline on average, up to 10.1%. This verifies that GATE ranks attribute values better than the baselines.

**Varying  $|\Sigma|$ .** Varying  $|\Sigma|$ , we evaluated the impact of the number of CCs in Figure 6(j). The accuracy of GATE, UncertainRule and Improve3C improves given more rules. For GATE, its  $F$ -measure changes from 0.805 to 0.866 when  $|\Sigma|$  varies from 20% to 100%. Indeed, the critic deduces more orders with more CCs for the creator to fine-tune its model, to get a higher accuracy in an earlier stage.

**Varying initial ts%.** We varied the ratio ts% of initial timestamps (randomly selected) from 4% to 20%. More initial timestamps help since (a) the creator has more training data and can have better

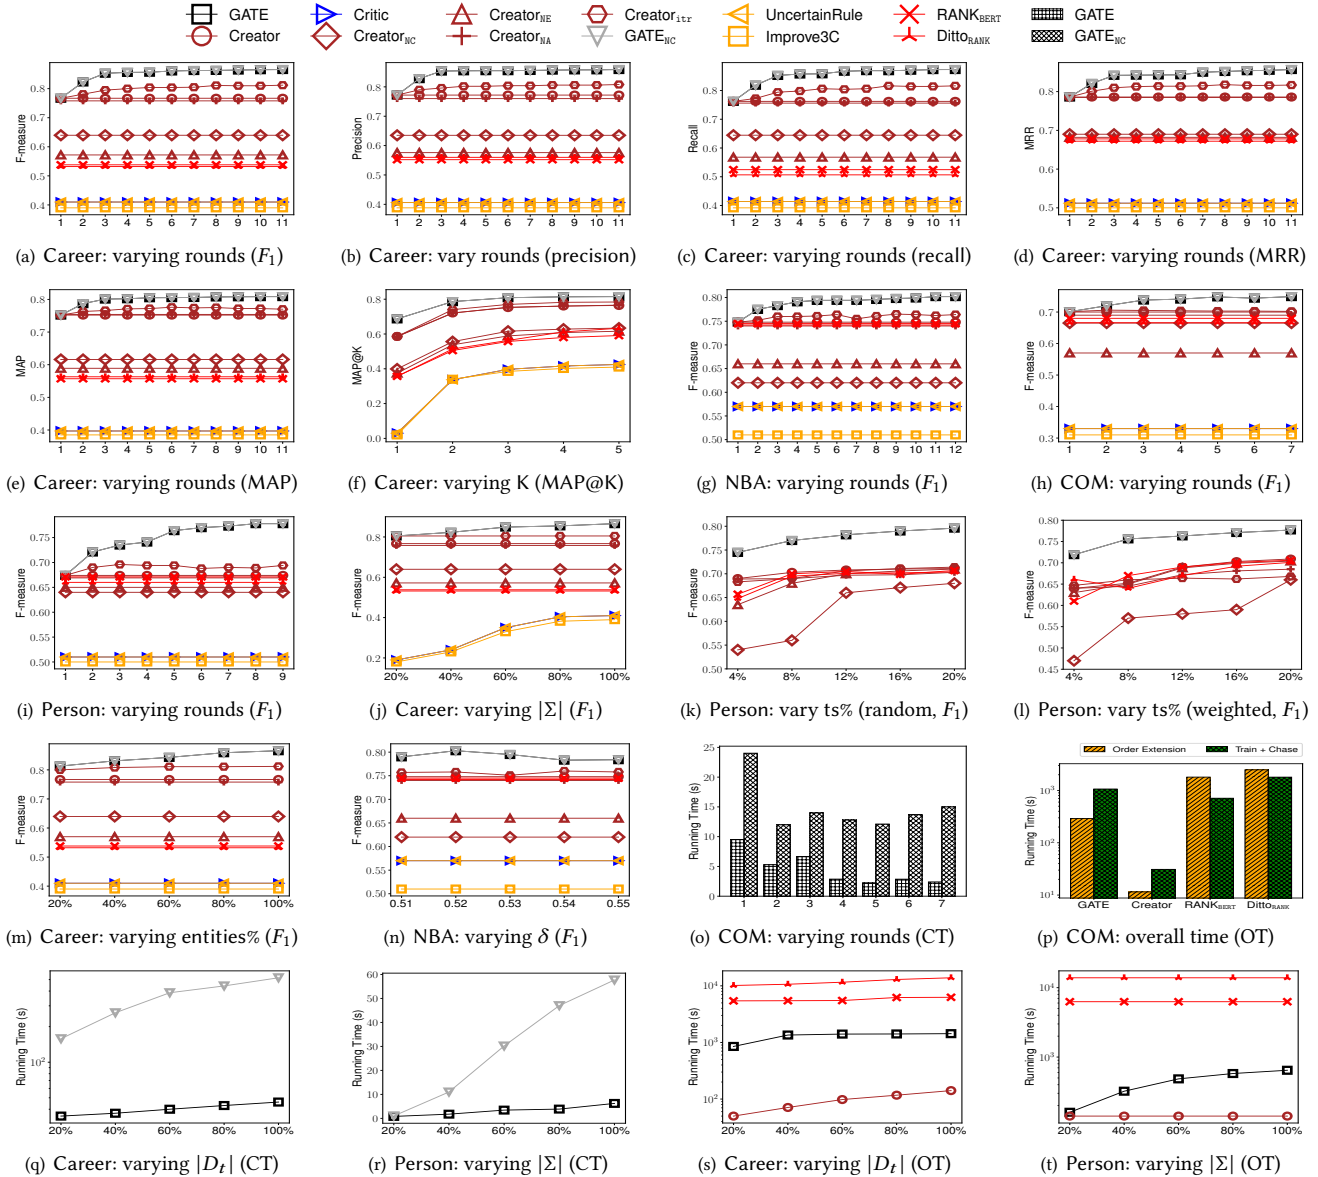


Figure 6: Performance evaluation

initial performance; and (b) the critic gets temporal orders as ground truth to perform deduction at the beginning of the chase. As shown in Figure 6(k), the  $F$ -measure of GATE increases (from 0.75 to 0.796 on Person) when ts% varies from 4% to 20%, as expected.

As remarked earlier, some attribute values may have more reliable timestamps than the others [40]. To study the impact of timestamp distributions, we assigned a weight to each attribute, where a larger weight indicates that the values of this attribute is more likely to be selected with initial timestamps, e.g., on Person, values on Status are more likely to have timestamps than Kids. Consistent with Figure 6(k), Figure 6(l) shows that the accuracy under weighted sampling is also improved with larger ts%. Note that the accuracy under weighted sampling (Figure 6(l)) is slightly lower than random sampling (Figure 6(k)), e.g., when ts%=20% on Person, its  $F$ -measure is 2.5% lower than random sampling. This is because weighted sampling is impacted by distribution discrepancy between training and

testing data, which is a common out-of-distribution issue for ML.

Varying entities%. As reported in Figure 6(m), we varied the percentage of entities that are used for the chase from 20% to 100%. As expected, GATE improves its  $F$ -measure since with more entities, the critic can deduce more orders via the chase, and the creator can have more augmented data to train the model, achieving higher accuracy. For instance,  $F$ -measure of GATE is improved by 5.3% on average.

Varying  $\delta$ . We next tested the impact of confidence threshold  $\delta$ . As shown in Figure 6(n), (a) although when  $\delta$  is small, less confident predictions may appear in subsequent deductions, more temporal orders could be used for training; (b) when  $\delta$  is too large, few ranked pairs learned are retained, and hence less augmented training data is returned. Since the creator receives less data to fine-tune its model, the accuracy may not be improved and it converges slowly. When  $\delta = 0.52$ , GATE has the highest accuracy on NBA. Thus, we

set  $\delta = 0.52$  as its default; for other datasets,  $\delta$  is set similarly.

**Exp-2: Efficiency.** We tested the efficiency of GATE, GATE<sub>NC</sub>, Creator, RANK<sub>Bert</sub> and Ditto<sub>Rank</sub>. Denoted by CT (resp. OT) the chase time (resp. the overall time, for training, chasing and order extension). For GATE, OT is accumulated over all rounds; its order extension time is the time for extending stable orders to total orders by  $\mathcal{M}_{\text{rank}}$  (Line 12 of Figure 3). For other baselines, the order extension time is the inference time of models for generating total orders. We did not report rule-based methods, which are fast, since they do not need to train models; but as shown in Exp-1, they are not accurate.

Chase time (rounds). We report CT of GATE and GATE<sub>NC</sub> in the iterative process. As shown in Figure 6(o), GATE is substantially faster than GATE<sub>NC</sub> for all rounds; it is 3.99X faster than GATE<sub>NC</sub> on average, up to 6.30X on COM. The speedup of GATE is due to the lazy evocation strategy we adopted, which accelerates the chase by maintaining designated structures. In contrast, GATE<sub>NC</sub> enumerates valuations and incurs redundant computation. Since GATE and GATE<sub>NC</sub> produces the same results, their total rounds are the same.

Overall time. We next report OT in Figure 6(p). Although GATE has multiple rounds, its overall time is comparable to most ML methods, e.g., GATE is 3.15X faster than Ditto<sub>Rank</sub> on average. In particular, the order extension time of GATE is smaller than most baselines except Creator, since GATE has to perform extra checks so that the total order generated is consistent with the known stable orders.

Varying  $|D_t|$ . We evaluated CT (resp. OT) of GATE and GATE<sub>NC</sub> (resp. ML methods) by varying  $|D_t|$  from 20% to 100% in Figure 6(q) (resp. 6(s)). With larger  $|D_t|$ , all methods take longer, as expected. Nonetheless, GATE is faster than GATE<sub>NC</sub> when  $|D_t|$  gets larger, since GATE maintains structures to avoid recomputation and does deduction pertaining to affected orders, e.g., GATE is 11.3X faster than GATE<sub>NC</sub> when  $|D_t|$  is 100%; the result for OT is consistent.

Varying  $|\Sigma|$ . We varied  $|\Sigma|$  from 20% to 100% on Person in Figure 6(r) and 6(t). As shown there, GATE is 7.54X faster than GATE<sub>NC</sub> on average, up to 12.2X, which again verifies the effectiveness of lazy evocation. OT of GATE increases as  $|\Sigma|$  is larger, since more training data (i.e., temporal orders) is deduced by CCs to train the ranking model.

**Exp-3: Case study.** We use Career to illustrate why GATE works.

(1) Initial prediction. In the first round, GATE trains the creator on data with initial timestamps. Due to the limited (5%) training data, its  $F$ -measure is 0.77 and some ranked pairs are mispredicted. One of confident and correct predictions is  $t_1 \leq_{\text{height}} t_2$ , where  $t_1$  and  $t_2$  denote the same player with heights 186cm and 188cm, respectively. While this player moved from team PARMA to SPAL, i.e.,  $t_1 \leq_{\text{team}} t_2$ , the creator makes a wrong prediction  $t_2 \leq_{\text{team}} t_1$ .

(2) Critic helps Creator. After the creator stage, the critic uses CCs to correct mispredicted temporal orders. By applying  $\varphi_7 : t_a \leq_{\text{height}} t_b \rightarrow t_a \leq_{\text{team}} t_b$  to the known  $t_1 \leq_{\text{height}} t_2$ , it deduces  $t_1 \leq_{\text{team}} t_2$ , correcting the mistake of Creator. Intuitively,  $\varphi_7$  holds since  $\leq_{\text{height}}$  is monotonic, and  $\leq_{\text{height}}$  and  $\leq_{\text{team}}$  correlate for young players.

Moreover, Critic provides augmented training data to Creator. For instance, if  $t_0 \leq_{\text{league\_name}} t_1$  and  $t_1 \leq_{\text{potential}} t_2$  are in the ground truth, the critic could apply CC  $\varphi_8 : t_a \leq_{\text{league\_name}} t_b \wedge t_b \leq_{\text{potential}} t_c \wedge t_a[\text{height}] \leq t_c[\text{height}] \rightarrow t_a \leq_{\text{position}} t_c$ , and

deduces a new pair  $t_0 \leq_{\text{position}} t_2$  that is unknown before. Here  $\varphi_8$  is learned from the data; intuitively, if a player moves to a new league (as indicated by monotonic  $\leq_{\text{height}}$ ) and if his potential changes, his position is likely adjusted, e.g., from LM to CAM. In the first 5 round, the critic creates 20K new ranked pairs per round as augmented training data, and the creator improves its model with the new data.

(3) Creator helps Critic. Critic cannot correctly deduce total orders from limited initial 5% timestamps. Nonetheless, with augmented training data provided by Critic, Creator can learn more ranked pairs with high confidence. On average it ranks hundreds of tuple pairs with high confidence in each round. These ranked pairs are in turn provided to Critic, for Critic to deduce more new ranked pairs.

(4) The creator learns better with more data. ML models are inclined to get more accurate when given more training data. Creator continually receives more augmented training data and incrementally trains its model accordingly. As a consequence, its  $F$ -measure increases from 0.77 to 0.82 (resp. 0.87) after the first (resp. last) round.

**Summary.** We find the following. (1) Combining deep learning and logic rules makes a promising approach to deducing currency. GATE is the most accurate, e.g., 0.866 in  $F$ -measure on Career, as opposed to 0.41 and 0.39 by rule-based UncertainRule and Improve3C, and 0.54 and 0.53 by ML-based RANK<sub>Bert</sub> and Ditto<sub>Rank</sub>. (2) GATE only takes 7 rounds to terminate on COM, which has 1,983,698 tuples. (3) On average, GATE is 34.4% and 7.8% more accurate than Critic and Creator, respectively, i.e., the creator and critic indeed benefit each other. (4) GATE beats GATE<sub>NC</sub> in efficiency (with the same accuracy) by 8.05X on average, up to 12.2X, verifying the usefulness of lazy evocation. (5) GATE has competitive overall time against existing ML methods, e.g., 1359s on COM, as opposed to 2514s by the fastest of them. Although rule-based methods are fast (they do not train models), their accuracy are low. (6) Creator (resp. GATE) is more accurate than its variants Creator<sub>NC</sub>, Creator<sub>NE</sub> and Creator<sub>NA</sub> (resp. Creator<sub>itr</sub>) by 8.1%, 10.9% and 1% (resp. 5.6%), respectively, verifying the need for context-aware embedding, chronological encoding and adaptive margin (resp. order justification), respectively.

## 7 CONCLUSION

The novelty of the work consists of the following. (1) We formulate a new problem for determining the timeliness of attribute values. (2) As a solution to the problem, we propose a creator-critic framework by combining deep learning and logic deduction, for the two to enhance each other. (3) We develop a novel ranking model to learn temporal orders on attribute values. (4) We show how to justify the learned orders, deduce more ranked pairs and provide feedback for the learner, by extending the chase using CCs. The experimental study has verified that GATE is promising in practice.

One future topic is to study how to catch conflicts and missing values given temporal orders. Another topic is to extend CCs [30] by embedding ranking models as predicates, to improve the ranking accuracy with logic conditions and interpret ranking in logic.

## ACKNOWLEDGMENTS

This work was supported in part by Royal Society Wolfson Research Merit Award WRM/R1/180014, Guangdong Basic and Applied Basic Research Foundation 2022A1515010120 and NSFC 62202313.

## REFERENCES

- [1] 2022. Full version. [https://github.com/yyssl88/Timeliness/blob/main/paper\\_full\\_version.pdf](https://github.com/yyssl88/Timeliness/blob/main/paper_full_version.pdf).
- [2] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Addison-Wesley.
- [3] Marcelo Arenas, Leopoldo Bertossi, and Jan Chomicki. 1999. Consistent Query Answers in Inconsistent Databases. In *PODS*.
- [4] Tobias Bleifuß, Sebastian Kruse, and Felix Naumann. 2017. Efficient Denial Constraint Discovery with Hydra. *PVLDB* 11, 3 (2017), 311–323.
- [5] Rpje Jagadeesh Bose, Rs Ronny Mans, and Van Der Wmp Wil Aalst. 2013. Wanna improve process mining results? It's high time we consider data quality issues seriously. In *Computational Intelligence & Data Mining*.
- [6] Philip Bramsen, Pawan Deshpande, Yoong Keok Lee, and Regina Barzilay. 2006. Inducing Temporal Graphs. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL.
- [7] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to rank using gradient descent. In *international conference on Machine learning*. 89–96.
- [8] Businesswire. 2022. Over 80 Percent of Companies Rely on Stale Data for Decision-Making. <https://www.businesswire.com/news/home/20220511005403/en/Over-80-Percent-of-Companies-Rely-on-Stale-Data-for-Decision-Making>.
- [9] Statistics Canada. 2022. Classification of legal marital status. <https://www23.statcan.gc.ca/imdb/p3VD.pl?Function=getVD&TVD=61748&CVD=61748&CLV=0&MLV=1&D=1>.
- [10] Nathanael Chambers and Dan Jurafsky. 2008. Jointly Combining Implicit Constraints Improves Temporal Ordering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Honolulu, Hawaii). ACL, 698–706.
- [11] Olivier Chapelle and Yi Chang. 2011. Yahoo! learning to rank challenge overview. In *Proceedings of the learning to rank challenge*. PMLR, 1–24.
- [12] Peter Christen and Ross W. Gayler. 2013. Adaptive Temporal Entity Resolution on Dynamic Databases. In *PAKDD*. Springer.
- [13] Xu Chu, Ihab F. Ilyas, and Paolo Papotti. 2013. Discovering Denial Constraints. *PVLDB* 6, 13 (2013), 1498–1509.
- [14] E. F. Codd. 1979. Extending the Database Relational Model to Capture More Meaning. *TODS* 4, 4 (1979), 397–434.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*. 4171–4186.
- [16] Ioannis Dikeoulas, Saadullah Amin, and Günter Neumann. 2022. Temporal Knowledge Graph Reasoning with Low-rank and Model-agnostic Representations. *CoRR* abs/2204.04783 (2022).
- [17] Xiaou Ding, Hongzhi Wang, Yitong Gao, Jianzhong Li, and Hong Gao. 2017. Efficient currency determination algorithms for dynamic data. *Tsinghua Science and Technology* 22, 3 (2017), 227–242.
- [18] Xiaou Ding, Hongzhi Wang, Jiakuan Su, Jianzhong Li, and Hong Gao. 2018. Improve3c: Data cleaning on consistency and completeness with currency. *arXiv preprint arXiv:1808.00024* (2018).
- [19] Xiaou Ding, Hongzhi Wang, Jiakuan Su, Muxian Wang, Jianzhong Li, and Hong Gao. 2020. Leveraging Currency for Repairing Inconsistent and Incomplete Data. *TKDE* (2020).
- [20] Aswathy Divakaran and Anuraj Mohan. 2020. Temporal Link Prediction: A Survey. *New Gener. Comput.* 38, 1 (2020), 213–258. <https://doi.org/10.1007/s00354-019-00065-z>
- [21] X. Dong, L. Berti-Equille, and D. Srivastava. 2009. Truth Discovery and Copying Detection in a Dynamic World. *PVLDB* 2, 1 (2009), 562–573.
- [22] X. L. Dong, L. Berti-Equille, Y. Hu, and D. Srivastava. 2010. Global detection of complex copying relationships between sources. In *PVLDB*.
- [23] Xuliang Duan, Bing Guo, Yan Shen, Yuncheng Shen, Xiangqian Dong, and Hong Zhang. 2020. Research on Parallel Data Currency Rule Algorithms. In *International Conference on Information Science and System*. 24–28.
- [24] Kevin K Duh. 2009. *Learning to rank with partially-labeled data*. University of Washington.
- [25] Exasol. 2020. Exasol Research Finds 58% of Organizations Make Decisions Based on Outdated Data. <https://www.exasol.com/news-exasol-research-finds-organizations-make-decisions-based-on-outdated-data/>.
- [26] Wenfei Fan, Floris Geerts, Xibei Jia, and Anastasios Kementsietsidis. 2008. Conditional functional dependencies for capturing data inconsistencies. *TODS* 33, 2 (2008), 6:1–6:48.
- [27] Wenfei Fan, Floris Geerts, Nan Tang, and Wenyuan Yu. 2013. Inferring data currency and consistency for conflict resolution. In *ICDE*. IEEE, 470–481.
- [28] Wenfei Fan, Floris Geerts, Nan Tang, and Wenyuan Yu. 2014. Conflict resolution with data currency and consistency. *Journal Data and Information Quality (JDIQ)* 5, 1-2 (2014), 6:1–6:37.
- [29] Wenfei Fan, Floris Geerts, and Jef Wijsen. 2011. Determining the currency of data. In *PODS*. ACM.
- [30] Wenfei Fan, Floris Geerts, and Jef Wijsen. 2012. Determining the Currency of Data. *TODS* 37, 4 (2012), 25:1–25:46.
- [31] Wenfei Fan, Ruochun Jin, Ping Lu, Chao Tian, and Ruiqi Xu. 2022. Towards Event Prediction in Temporal Graphs. *PVLDB* 15, 9 (2022), 1861–1874.
- [32] Wenfei Fan, Ping Lu, and Chao Tian. 2020. Unifying Logic Rules and Machine Learning for Entity Enhancing. *Sci. China Inf. Sci.* 63, 7 (2020).
- [33] Wenfei Fan, Chao Tian, Yanghao Wang, and Qiang Yin. 2021. Parallel Discrepancy Detection and Incremental Detection. *PVLDB* 14, 8 (2021), 1351–1364.
- [34] Shenzhen Municipal Government. 2022. Self-employed Entrepreneurs. [https://opendata.sz.gov.cn/data/dataSet/toDataDetails/29200\\_01300931](https://opendata.sz.gov.cn/data/dataSet/toDataDetails/29200_01300931).
- [35] Tanya Goyal and Greg Durrett. 2019. Embedding Time Expressions for Deep Temporal Ordering Models. In *Conference of the Association for Computational Linguistics (ACL)*. ACL.
- [36] Shuang Han, Xuanhui Wang, Mike Bendersky, and Marc Najork. 2020. Learning-to-Rank with BERT in TF-Ranking. *arXiv preprint arXiv:2004.08476* (2020).
- [37] Yujing Hu, Qing Da, Anxiang Zeng, Yang Yu, and Yinghui Xu. 2018. Reinforcement learning to rank in e-commerce search engine: Formalization, analysis, and application. In *SIGKDD*. 368–377.
- [38] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- [39] James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. Overcoming catastrophic forgetting in neural networks. *CoRR* abs/1612.00796 (2016).
- [40] Angelina Prima Kurniati, Eric Rojas, David Hogg, Geoff Hall, and Owen A Johnson. 2019. The assessment of data quality issues for process mining in healthcare using Medical Information Mart for Intensive Care III, a freely available e-health record database. *Health informatics journal* 25, 4 (2019), 1878–1893.
- [41] Stefano Leone. 2022. FIFA 22 complete player dataset. <https://www.kaggle.com/stefanoleone992/fifa-21-complete-player-dataset>.
- [42] Furong Li, Mong-Li Lee, Wynne Hsu, and Wang-Chiew Tan. 2015. Linking Temporal Records for Profiling Entities. In *SIGMOD*. ACM, 593–605.
- [43] Mohan Li and Jianzhong Li. 2016. A minimized-rule based approach for improving data currency. *J. Comb. Optim.* (2016), 812–841.
- [44] Mohan Li, Jianzhong Li, Siyao Cheng, and Yanbin Sun. 2018. Uncertain rule based method for determining data currency. *IEICE TRANSACTIONS on Information and Systems* 101, 10 (2018), 2447–2457.
- [45] Mohan Li and Yanbin Sun. 2018. Currency Preserving Query: Selecting the Newest Values from Multiple Tables. *IEICE TRANSACTIONS on Information and Systems* 101, 12 (2018), 3059–3072.
- [46] Pei Li, Xin Luna Dong, Andrea Maurino, and Divesh Srivastava. 2011. Linking Temporal Records. *PVLDB* 4, 11 (2011), 956–967.
- [47] Yuliang Li, Jinfeng Li, Yoshihiko Suhara, AnHai Doan, and Wang-Chiew Tan. 2020. Deep Entity Matching with Pre-Trained Language Models. *PVLDB* 14, 1 (2020), 50–60.
- [48] Yu Liang, Xuliang Duan, Yuanjun Ding, Xifeng Kou, and Jingcheng Huang. 2019. Data Mining of Students' Course Selection Based on Currency Rules and Decision Tree. In *International Conference on Big Data and Computing*. 247–252.
- [49] Ashley Little. 2020. Outdated Data: Worse Than No Data? <https://info.aldensys.com/joint-use/outdated-data-is-worse-than-no-data#:~:text=Obsolete%20data%20about%20the%20condition,too%20old%20to%20be%20reliable>.
- [50] Tie-Yan Liu. 2009. Learning to Rank for Information Retrieval. *Found. Trends Inf. Retr.* 3, 3 (2009), 225–331. <https://doi.org/10.1561/1500000016>
- [51] Tie-Yan Liu. 2010. Learning to rank for information retrieval. In *SIGIR*.
- [52] Ester Livshits, Alireza Heidari, Ihab F. Ilyas, and Benny Kimelfeld. 2020. Approximate Denial Constraints. *PVLDB* 13, 10 (2020), 1682–1695.
- [53] Niels Martin, Antonio Martínez-Millana, Bernardo Valdivieso, and Carlos Fernandez-Llatas. 2019. Interactive Data Cleaning for Process Mining: A Case Study of an Outpatient Clinic's Appointment System. In *International Conference on Business Process Management*.
- [54] Qiang Ning, Zhili Feng, and Dan Roth. 2017. A Structured Learning Approach to Temporal Relation Extraction. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 1027–1037.
- [55] Qiang Ning, Hao Wu, Haoruo Peng, and Dan Roth. 2018. Improving Temporal Relation Extraction with a Globally Acquired Statistical Resource. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. ACL, 841–851.
- [56] Rodrigo Frassetto Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *CoRR* abs/1901.04085 (2019).
- [57] Rama Kumar Pasumarthi, Sebastian Bruch, Xuanhui Wang, Cheng Li, Michael Bendersky, Marc Najork, Jan Pfeifer, Nadav Golbandi, Rohan Anil, and Stephan Wolf. 2019. TF-ranking: Scalable tensorflow library for learning-to-rank. In *SIGKDD*. 2970–2978.
- [58] Eduardo H. M. Pena, Eduardo Cunha de Almeida, and Felix Naumann. 2019. Discovery of Approximate (and Exact) Denial Constraints. *PVLDB* 13, 3 (2019), 266–278.

- [59] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*.
- [60] Royal Mail. 2018. Dynamic Customer Data in a Digital World: Data Services Insight Report. <https://www.royalmail.com/business/system/files/royal-mail-data-services-insight-report-2018.pdf>.
- [61] Ali Sadeghian, Mohammadreza Armandpour, Anthony Colas, and Daisy Zhe Wang. 2021. ChronoR: Rotation Based Temporal Knowledge Graph Embedding. In *AAAI*. AAAI Press, 6471–6479.
- [62] Fereidoon Sadri and Jeffrey D. Ullman. 1980. The Interaction between Functional Dependencies and Template Dependencies. In *SIGMOD*.
- [63] Yi Tay, Minh C Phan, Luu Anh Tuan, and Siu Cheung Hui. 2017. Learning to rank question answer pairs with holographic dual lstm architecture. In *SIGIR*. 695–704.
- [64] Julien Tourille, Olivier Ferret, Aurélie Névéol, and Xavier Tannier. 2017. Neural Architecture for Temporal Relation Extraction: A Bi-LSTM Approach for Detecting Narrative Containers. In *ACL*. ACL, 224–230.
- [65] Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs. In *International Conference on Machine Learning (ICML)*, Vol. 70. PMLR, 3462–3471.
- [66] Hongzhi Wang, Xiaoou Ding, Jianzhong Li, and Hong Gao. 2018. Rule-based entity resolution on database with hidden temporal information. *TKDE* 30, 11 (2018), 2199–2212.
- [67] Jun Xu, Xiangnan He, and Hang Li. 2020. Deep Learning for Matching in Search and Recommendation. *Found. Trends Inf. Retr.* 14, 2-3 (2020), 102–288.
- [68] Jing Yao, Zhicheng Dou, Jun Xu, and Ji-Rong Wen. 2021. RLPS: A Reinforcement Learning-Based Framework for Personalized Search. *TOIS* 39, 3 (2021), 1–29.
- [69] Jingran Zhang, Fumin Shen, Xing Xu, and Heng Tao Shen. 2020. Temporal Reasoning Graph for Activity Recognition. *IEEE Trans. Image Process.* (2020).
- [70] Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2016. Listwise ranking functions for statistical machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24, 8 (2016), 1464–1472.