

# Creative influence prediction using graph theory

Francesco Alfieri<sup>2</sup>, Luigi Asprino<sup>1</sup>, Nicolas Lazzari<sup>1,†</sup> and Valentina Presutti<sup>1</sup>

<sup>1</sup>Department of Computer Science and Engineering, University of Bologna, Mura Anteo Zamboni, 7, Bologna 40126, Italy

<sup>2</sup>LILEC, University of Bologna, Via Cartoleria, 5, Bologna 40124, Italy

## Abstract

Creative influence is responsible for a considerable part of the creative process of an artist and can largely be associated with their social circle. It has been observed that the type and amount of relationships with other fellow artists correlates with the success of an artist. Most of the recent literature has focused on using artefact similarity as a proxy for creative influence between two artists. However, this approach neglects the significance of an artist's social network or flattens the individuality of a relationship by only addressing it as a direct connection. In this work, we propose an ontology to comprehensively model the relationship between individuals as a Knowledge Graph. Additionally, we design and implement an explainable method based on graph theory to predict the influences of an artist given their social network. We evaluate our method on a dataset of relationships between Jazz musicians and achieve accurate results when compared to baselines that rely on the distribution of the data. Our results are aligned with relevant works from the socio-cognitive and psychology fields. <sup>1</sup>

## Keywords

computational creativity, graph theory, knowledge graph, artistic influence

## 1. Introduction

Identifying, capturing and hypothesising the influences of an artist is an important aspect that an art critic considers when analysing an artefact or, in general, the artist themselves [1]. The main difficulty in determining influence lies in the subjective nature of the problem. Identifying the influence of an artist on another artist requires a profound knowledge of both entities, their geographical location, the socio-cultural context in which they lived, the technicalities of their artefacts, and so on. It has been argued that an unambiguous definition of creative influence is problematic [2].

Yet, the importance of capturing and understanding the influences of an artist has a great impact on many pragmatic aspects. Mitali and Ingram [3] analyses the work of 90 pioneers in the abstract art movement. The results provide very strong evidence that the success and fame of an artist are related to their social relationships. While it is true that creativity fosters

---

<sup>1</sup>The code and the ontology developed is shared at [https://github.com/n28div/influence\\_prediction](https://github.com/n28div/influence_prediction) under CC-BY license.

CREAI 2023 - Workshop on Artificial Intelligence and Creativity, November 6-9, 2023, Rome, Italy

<sup>†</sup> Corresponding author.

✉ francesco.alfieri5@studio.unibo.it (F. Alfieri); luigi.asprino@unibo.it (L. Asprino); nicolas.lazzari3@unibo.it (N. Lazzari); valentina.presutti@unibo.it (V. Presutti)

🆔 0009-0007-2293-5682 (F. Alfieri); 0000-0003-1907-0677 (L. Asprino); 0000-0002-1601-7689 (N. Lazzari); 0000-0002-9380-5160 (V. Presutti)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

those relationships, the more an artist gets deeper into a clique formed by other meaningful artists, the more its work is acclaimed by critics. For instance, the success of the band *The Velvet Underground* is often partially associated to their relationship with the artist *Andy Warhol* [4]. While the study focuses on visual arts, the same argument holds for any other artistic endeavour. In music, the importance of the relationships of a musician in its creative process is widely recognised [5, 6]. A popular example is the *teacher-student* relationship, sometimes referred to as *mentor-pupil*. Famous artists are often the students of other famous artists, which results in an inevitable influence on their creative process [7].

To understand the influence on an artist, it is important to take into account its social relationships as well. Most recent approaches, however, use artefact similarity as a proxy for creative influence Abe et al. [8], Saleh et al. [9], Elgammal and Saleh [10], O’Toole and Horvát [11], Park et al. [12]. While this has resulted in promising outcomes, it neglects the incidence of the social circle, making it impossible to detect influences from artists whose stylistic genres are different. Moreover, by persuading a solely similarity-based approach, it is unfeasible to detect artistic influences between two artists that perform on different domains, such as influences of painters on musicians [13, 4], such as in the example of *The Velvet Underground*.

In this work, we propose a method to predict the influence of musical artists by only taking into account their social network. We design an ontology to model the relationships between artists in an expressive way. Rather than consider them as simple direct relationships we model them as complex situations that involve different agents and concepts. We refactor the data from the Linked Jazz project [14]<sup>1</sup>, a Knowledge Graph encoding curated relationship between Jazz musicians (among which creative influence), to comply with our ontology and rely on it as a ground truth to perform influence prediction. We frame influence prediction as a classification task where one has to identify and rank artists according to their likelihood to be influential for a given artist. Our approach is based on techniques from graph theory, namely the  $f$ -communicability of a graph [15]. Informally, the  $f$ -communicability of a graph provides information on how close two artists are as a function of their connections. The shorter the connections between two artists, the higher their communicability. We consider the  $f$ -communicability between nodes  $i$  and  $j$  as the degree of influence that  $j$  has on  $i$  or, in other words, how influenced is  $i$  with respect to  $j$ . Our method assigns a weight to each relation in the Knowledge Graph based on the type of relationship. We approximate the weighting function by maximising the  $f$ -communicability between influential relationships asserted in the original Knowledge Graph in an optimisation procedure. We evaluate our results through the use of standard information retrieval measures (MRR, MAP, DCG) and compare our method with baselines that rely on the distribution of the data. The learned weighting function obtains results that are aligned with other relevant studies from the socio-cognitive and psychology fields.

Our contributions can be summarised as follows:

- an ontology to model the relationships between human agents, with a particular focus on artists;
- an explainable method to predict creative influence between artists.

---

<sup>1</sup>Retrieved from <https://triplifydb.com/pratt/linked-jazz/>

The paper is structured as follows: in Section 2 we provide a review of related works addressing the prediction and identification of influence between artists. In Section 3 we present the ontology and its associated method to predict creative influence between artists. In Section 4 we describe the experimental setup while in Section 5 we present the obtained results. Section 6 summarises the results of previous sections and highlights potential extensions and future work.

## 2. Related Works

Evaluating the influences of an artist, and in particular of a composer or a musician, is mostly considered a subjective task. Usually, experts analyse the compositions of an artist in a critical way to relate them to other important artists. One approach to detecting creative influence is to directly analyse explicit influence connections, curated by human annotators. Smith and Georges [16], for example, analyses the influences identified in the Classical Music Navigator (CMS) to better understand the influence of the composers in the dataset. A similar approach is taken by Georges and Seckin [17], where the data on creative influence is used to investigate the similarity of musical compositions.

Relying on similarity as a proxy for creative influence is a popular approach that has been explored using different techniques. Abe et al. [8] define a framework where influence can be modelled using a graph structure. Edges are added to the graph by taking into account the similarity between the two artworks. Several works have explored this approach in the visual art domain with promising results [9, 10], and in the musical domain. O’Toole and Horvát [11] models the influence of musical composition as the probability of success of a composition given its similarity to other popular compositions. In Park et al. [12], the influence of a composer on another composer is measured as the degree of similarity between their compositions. A composer is classified as influential when musical features of its compositions are re-used by other composers.

Relying on artefact similarity, however, can be a problematic approach in art. Influence can affect an artist in a negative way, in the sense that the influenced artist deliberately abstains from his influence [2]. These kinds of artists are sometimes defined as deviant artists [18, 19]. Mauskapf et al. [20] investigates similarity with respect to socio-cultural indicators, such as geographical and temporal location or organisational system in which the artist lives. Findings suggest that highly embedded individuals, i.e. individuals with a dense social network, are more likely to produce novel artefacts that can be influential to other artists. Borowiecki [21] analyses influence of the *teacher-student* relationship through a combination of artifact features. Albeit with different intensities, findings highlight the importance of such a relationship, as also observed in Simonton [7]. Analysing the social network of artists using complex network tools has been explored in literature [10, 22, 23]. The relationships taken into account are often the result of heuristic methods or are limited to a few relationship types, such as *teacher-student* or *bandmates*. Relying on a rich social network where different relationship types are taken into account has been proven to be an effective way of uncovering meaningful insights [24] from data. Moreover, considering many different relationship types is an important requirement, as it has been largely discussed how different relationship ties can contribute differently to creative

**Table 1**

---

Relationship types in the ontology
Admiration, Fellowship, Bandmate, Copupils, Friendship, Mentorship, Parenthood, Rivalship, Siblinghood

---

influence [7, 23, 25].

Differently from the described approaches, we investigate the importance of social relationships without taking into account any information on the creative artefact. Our approach can be easily integrated with other methods that use similarity, to yield a more general method for uncovering hidden relations when perceptual similarity is the only measure taken into account.

### 3. Methodology

This section provides a detailed description of our method. In Section 3.1 we discuss the design and implementation of the ontology. In Section 3.2 we describe in detail the algorithm used to compute the influences between entities in the KG. In Section 3.3 we describe the procedure designed to learn the weight of each relationship type.

#### 3.1. Relationship Ontology

The ontology is built upon the concept of *social relation* from the DOLCE ontology [26]. A *social relation* can be defined as a situation expressed by some source of information<sup>2</sup>, some participants and a role that qualifies the type of the relation.

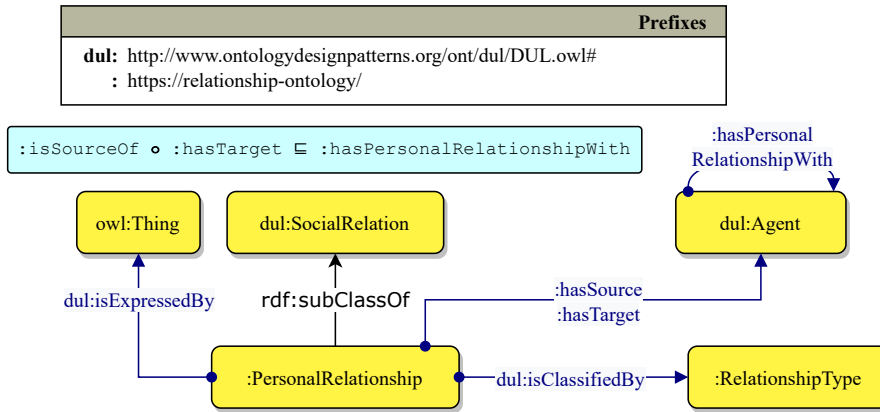
In our ontology, we only take into account pairwise relationships. This is intended as two roles (*source* and *target*) that partake in the relationship. In this way, it is possible to define a relation between sets of entities while retaining the directedness of the relationship. An example is the *Mentorship* relation, where a single mentor might have multiple students. The source of information tracks the provenance of the relationship assertion. The relationship types are described in Table 1.

Figure 1 visually represents the ontology. The class `:PersonalRelationship` reifies the relationship between two agents. To guarantee compactness and easiness of querying, we define the property `:hasPersonalRelationshipWith` to instantiate the binary projection reified by `:PersonalRelationship`. This is defined through the use of property chain axioms (blue box in Figure 1). Note that we do not restrict the amount of subjects to the `:hasSource` and `:hasTarget` properties. This enables the representation of pairwise relationships between two sets of entities.

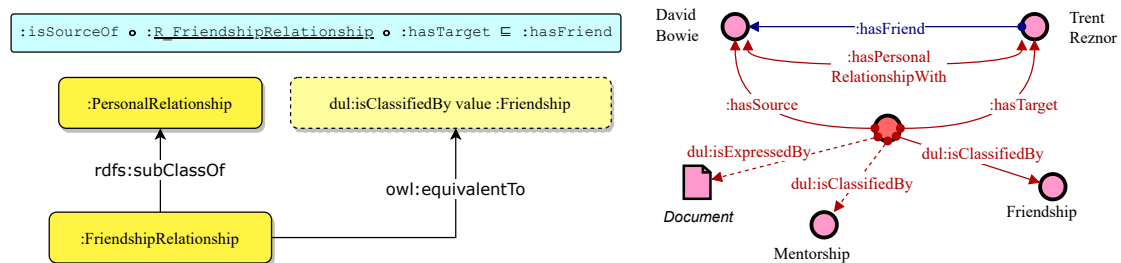
In Figure 2 an example of how the ontology can be used to define the friendship relationship of Table 1 is described. In order to define a friendship relationship between the musicians *Trent Reznor* and *David Bowie* it is sufficient to add the triple `< Trent Reznor, :hasFriend, David Bowie >` in the Knowledge Graph. The reification of the relationship is automatically performed by the inference engine.

---

<sup>2</sup>An information object in DOLCE



**Figure 1:** Ontology in Grafoo syntax. A pairwise relationship involving two agents is reified as a `PersonalRelationship` classified by an arbitrary type and expressed by another entity. A binary projection is automatically inferred using the property chain axiom in the blue box.



**Figure 2:** Example of the friendship relationship using the ontology. The `FriendshipRelationship` class (left) is constrained to be described by the `Friendship` class. The relationship  $\langle \text{Trent Reznor}, :hasFriend, \text{David Bowie} \rangle$  is represented on the right. The red colour (bottom right) identifies the inferred axioms and entities. Dashed arrows are refinement additions. The property `:R_FriendshipRelationship` (top left) is used to express class membership in property chain axioms using the rolification technique [27].

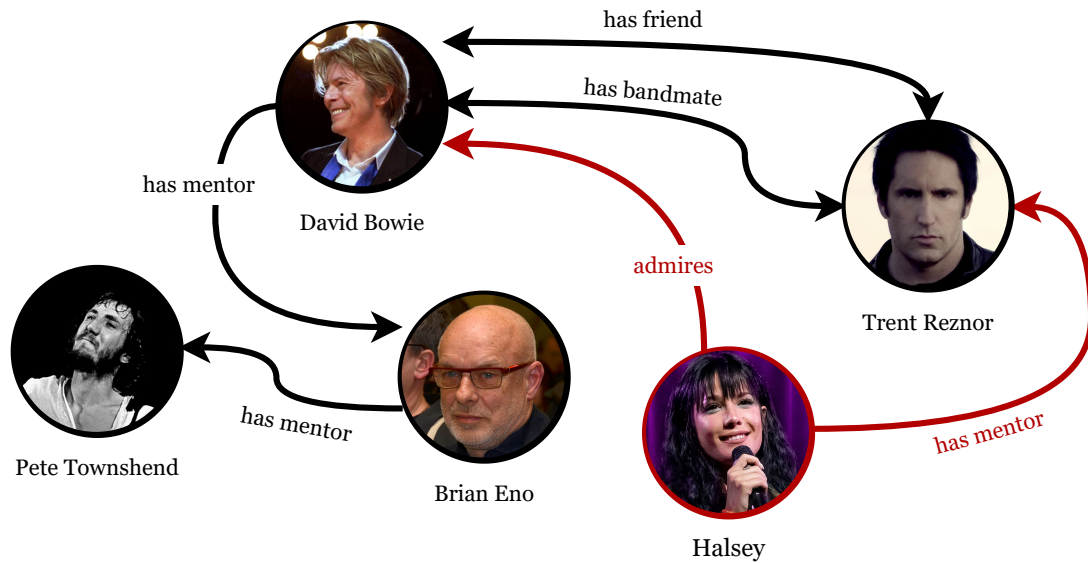
The implemented reification allows us to represent the relationship as a whole rather than flattening it into a binary relation, resulting in a richer characterisation of the relationship and a high degree of control in further refining it. For example, in Figure 2 the dashed properties represent refinement operations over the initial definition. It is possible to classify the relationship as both a friendship and mentorship relationship while adding documents that act as references to back up the assertion.

### 3.2. *f*-communicability as influence indicator

Once relations are represented using the ontology described in Section 3.1, the resulting Knowledge Graph can be interpreted as a semantically defined social network. By only relying on the binary projections of the reified relationships we can extract a directed graph  $G$  where entities are directly connected to each other by means of a set of edges  $E$ . In order to quantify

the influence of one artist on another, we exploit tools that pertain to the analysis of complex networks.

Our approach is based on the *f-communicability* [15] of the nodes in a graph, which is defined as a function of the paths that connect two distinct nodes. Generally, a node is *highly communicative* with another node if there are many paths that connect the two nodes. The length of the connecting path is an important factor that needs to be considered. If information (e.g. creative influence) has to travel in the graph, a shorter path is much more convenient than a longer one. This means that artists are generally influenced by close connections. Nonetheless, long connections should not be ignored. The *f-communicability* score between two connected nodes changes as a function of the length of the path: the shorter the path, the higher the communicability of the nodes.



**Figure 3:** Example of social graph. Intuitively, the communicability between *Halsey* and *David Bowie* should be high given the direct and non-direct connections. Nonetheless, the communicability with *Pete Townshend* should be considered as well, since there is a chain of relations that connects the two artists.

In Figure 3 a visual example of the communicability between artists and the importance of the length of the path is reported.

We rely on *f-communicability* as an indicator of how influential an artist is with respect to the other artists in the Knowledge Graph. We define a Knowledge Graph  $KG$  as a directed edge-labelled graph  $G = (V, E, L)$  where  $V$  represents the set of nodes in the graph,  $E$  the set of edges and  $L$  the set of labels that can be assigned to an edge  $e \in E$ .  $L$  is effectively the set of binary projections obtained from the relationship types of Table 1.

The *f-communicability* between nodes  $i, j$  in a graph  $G$  is computed as  $f(A)_{ij}$  where  $f$  is a suitable matrix function and  $A$  is the adjacency matrix of  $G$ . In order to take into account the different importance of the relationship types  $l \in L$  we use a weighting function  $w : L \rightarrow \mathbb{R}_0^+$ . The function  $w$  can be interpreted as the absolute degree of importance of a relationship type

$$\begin{array}{c}
\begin{array}{ccccc}
& PT & BE & DB & TR & H \\
PT & \left( \begin{array}{ccccc}
0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 2 & 0 \\
0 & 0 & 2 & 0 & 0 \\
0 & 0 & 0 & 1 & 0
\end{array} \right) & & & & \\
BE & & & & & \\
DB & & & & & \\
TR & & & & & \\
H & & & & & 
\end{array} & 
\begin{array}{ccccc}
\left( \begin{array}{ccccc}
0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 \\
0.5 & 1 & 2 & 2 & 0 \\
0 & 1 & 2 & 2 & 0 \\
0 & 0 & 1 & 1 & 0
\end{array} \right) & & & & \\
& & & & & \\
& & & & & \\
& & & & & \\
& & & & & \\
& & & & & 
\end{array} & 
\begin{array}{ccccc}
\left( \begin{array}{ccccc}
0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 \\
0.5 & 1.67 & 2 & 3.33 & 0 \\
0.33 & 1 & 3.33 & 2 & 0 \\
0 & 0.33 & 1 & 1.67 & 0
\end{array} \right) & & & & \\
& & & & & \\
& & & & & \\
& & & & & \\
& & & & & \\
& & & & & 
\end{array}
\end{array}
\end{array}$$

(a)  $f^1(A_{ij})$                       (b)  $f^2(A_{ij})$                       (c)  $f^3(A_{ij})$

**Figure 4:** Examples of equation 3 computed on the example from Figure 3. The labels  $TD$ ,  $DB$ ,  $BE$ ,  $TR$  and  $H$  are used to represent the nodes of, respectively, Pete Townshend, Brian Eno, David Bowie, Trent Reznor and Halsey. For simplicity of understanding the weights of equation 1 are considered to be  $w(l) = 1 \quad \forall l \in L$ . Relations that represent direct influence (e.g.  $H \rightarrow DB$ ) are not added to the graph. In Figure a the adjacency matrix obtained from Equation 1 is shown. Note that the nodes that have redundant connections (e.g.  $DB \rightarrow TR$ ) have a value of 2. In Figure b the  $f$ -communicability function is computed for walks of length 1. Nodes that were previously disconnected (e.g.  $TR \rightarrow PT$ ) are now connected. The same happens in Figure c, except that new connections have a lower weight (e.g.  $H \rightarrow PT$ ).

and can either be predefined or learned, as we show in section 3.3. We can now define the weighted adjacency matrix  $A$  as

$$A_{ij} = \sum_{l \in \text{labels}(i,j)} w(l) \quad (1)$$

where  $\text{labels}(i, j)$  is the set of labels of the edges between  $i$  and  $j$ .

The  $f$ -communicability  $f(i, j)$  between the nodes  $i, j \in V$  is computed as

$$f(i, j) = \sum_{k=1}^{\infty} \alpha_k (A^k)_{ij} \quad (2)$$

where  $A^n$  is the  $n$ -th power of  $A$ , and  $\alpha_k$  is a weight assigned to the walks of length  $k$ . Given an adjacency matrix  $A$  of a graph  $G$ , the entry  $(A^k)_{ij}$  is equal to the sum of the weights of all walks in  $G$  from node  $i$  to node  $j$  of length exactly  $k$  [28].

The weight  $\alpha_k$  defines how the importance of a walk should decay as a function of its length and needs to be carefully chosen in order to make sure that the sum of Equation 2 converges to the finite value. This is done by using a succession  $(\alpha_k)$  converging to 0 [15]. This ensures that walks of length  $\infty$  will have a null weight. The choice of how  $\alpha_k$  should decay leads to the use of different functions [29, 30]. We follow the definition of Estrada and Hatano [30] and set  $\alpha_k = \frac{1}{k!}$ . Our function  $f$  is hence the exponential matrix function. Given the computational cost of obtaining the exact exponential function (particularly for large graphs), we define our actual  $f$ -communicability function as an approximation of the exponential matrix function. This is done by truncating the power series that represents such a function. Formally, we compute

$$f^D(i, j) = \sum_{k=1}^D \frac{(A^k)_{ij}}{k!} \quad (3)$$

where  $D$  is a parameter corresponding to the maximum length of a walk taken into account by  $f^D$ . An approximation of the centrality of a node  $i \in G$  can be obtained by computing  $f^D(i, i)$ . See Figure 4 for an example on how Equation 3 is computed based on Figure 3.

### 3.3. Learning the importance of a relationship

One requirement of the method just described is that the graph from which the influence is predicted is weighted.

Defining the weight of a social relationship is an elaborate task, particularly when it needs to be contextualised in the creative domain. Perry-Smith [23] argues on the existence of strong and weak social ties and their influence on creativity. Strong ties, defined as highly redundant connections between two individuals (e.g. *Trent Reznor* and *David Bowie* in Figure 3) are found to positively correlate with creativity. A straightforward approach, which will serve as a baseline, is to assign to each relation type the same weight (e.g. 1) by using  $w : L \rightarrow 1$ . However, it is important to note that not all relationships equally correlate with creativity [7]. Given a reference KG, a distributional approach can be taken by defining the weight of the label to be (inversely) proportional to their distribution in the graph. This can be done by defining  $w(l) = \frac{|E_l|}{|E|}$  (respectively  $w(l) = \frac{|E_l|}{|E_l|}$ ) where  $E_l = \{e \in E \text{ s.t. } l \in \text{labels}(e)\}$ . This approach assumes that the importance of a relationship is (inversely) proportional to the distribution of that same relationship among the reference population. While this might be true, KG relies on the open-world assumption, where data incompleteness is taken into account. This is an important aspect since it is difficult (if not impossible) to completely enumerate the social relationships of an individual.

We propose to learn the weights assigned by  $w$  by fitting the data in the knowledge graph. This can be obtained by framing the problem of predicting creative influence as a multi-class classification problem. Given an edge  $e = (i, j)$  we can interpret  $f^D(i, j)$  as the probability that  $e \in E$  with  $l \in \text{labels}(e)$  where  $l$  is the label assigned to the edges that semantically states that  $i$  is creatively influenced by  $j$ . Essentially, the  $f$ -communicability of a pair of nodes  $(i, j)$  measures the degree to which  $i$  is creatively influenced by  $j$ .

In order to do that we minimise the cross entropy-based learning-to-rank loss defined in Bruch [31] between  $f^D(\hat{A}_{ij})$  and  $A_{tij}$  where  $t$  is the label assigned to the relationship that represents the influence of an artist onto another artist and  $\hat{A}_{ij} = \sum_{l \in L_t} A_{lij}$ . By relying on a learning-to-rank loss we learn weights in such a way that, given an input artist, its most influential artists are given a high weight. As a result, the model will be able to rank other artists in a meaningful way despite the absence of any explicit edge in the original graph.

While aggregating relations as in Equation 1 is a natural and straightforward approach, it is reasonable to suppose that the joint presence of two relationships, e.g. *friendship* and *mentorship* together, might be more (or less) important than the sum of the two components. To take this additional consideration into account we perform a non-linear combination using a feedforward neural network with one hidden layer using a ReLU activation. Equation 1 is hence updated to

$$A_{ij} = NN([A_{tij} \quad \forall t \in L]) \quad (4)$$

with  $NN$  being the function learned by the neural network.



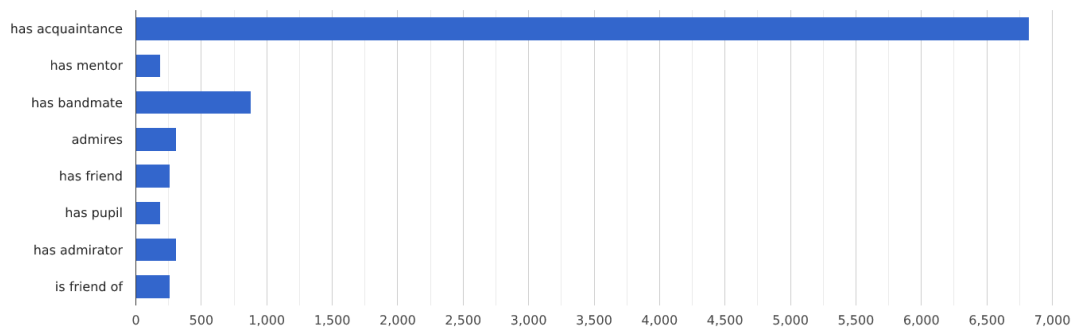
## 4. Experiments

In this section, we describe the experiment that we perform on the method proposed in Section 3. In Section 4.1 we describe the creation of the Knowledge Graph that is used to estimate and evaluate the predicted creative influence while in Section 4.2 we describe the experimental setup used to assess the accuracy of the method from Section 3.

### 4.1. Knowledge Graph

We rely on the data from the Linked Jazz project [14]. Linked Jazz is a Knowledge Graph containing information about famous jazz musicians and their social connections to other musicians. Data is semi-automatically annotated from the transcription of artists' interviews using crowd-sourced annotations. While some relationship types are objective (e.g. *bandmate* relationship) some have a subjective definition (e.g. *influence* relationship) and needs to be interpreted in the context of the interview. Annotators are provided with a definition for each relationship type, which partly addresses this issues. Modelling social relations as linked open data has shown how it is possible to uncover meaningful relationships between entities that are otherwise difficult to uncover [24].

We align the Linked Jazz KG to our ontology (described in Section 3.1) through the use of a SPARQL construct query.



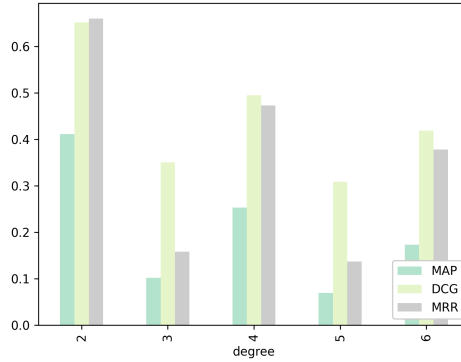
**Figure 5:** Relation distribution in the KG.

The KG contains a total of 5058 statements between 70 artists, where each musician has 72 relations asserted on average. In Figure 5 the distribution of the relationship between the entities in the KG is shown.

### 4.2. Experimental setup

We experiment with the methods of Section 3 on the Knowledge Graph described in the previous section.

In order to learn the weights of the function  $w$  of Equation 1 we split the data into the usual training and testing partitions, where the testing partition is 20% of the total data. The resulting training and testing data are hence composed of, respectively, 56 and 14 artists. Since the



**Figure 6:** Aggregate value combining all methods as a function of the variable  $D$  in Equation 2.

amount of explicitly stated influences available in the KG is much lower than the total amount of edges, the split is only based on nodes that have some influence edges asserted. This allows us to effectively evaluate the accuracy of the model with respect to creative influence prediction. Given the small size of both splits, we evaluate the models on aggregate values from a set of 5 distinct experiments on different subsets of the data. This helps us mitigate the noise due to the low amount of data available.

We evaluate each model using Mean Reciprocal Rank (MRR), Mean Average Precision (MAP) and Discounted Cumulative Gain (DCG). All the listed metrics measure how high are ranked appropriate values, e.g. how high are ranked actual influential artists with respect to a reference artist. MRR can be interpreted as how far is the first influential artist in the ordered list. MAP is the average of the number of relevant entries within the first  $k$  results, where  $k$  is the number of influences of each artist. DCG evaluates the results by penalising when relevant entries are not positioned at the top of the list. Each model is trained using COCOB [32], a parameter-free optimisation method. All the experiments are performed on an Intel i9 with 128GB of RAM and an Nvidia RTX3090 with 24GB of VRAM.

## 5. Results

In Figure 6 results from all methods are aggregated and plotted as a function of the degree  $D$  of Equation 2. The results highlight how using walks whose length extends at most up to 2 nodes obtains the best results. Influences from artists that are difficult to reach, i.e. that are separated by many other nodes, add noise to the  $f$ -communicability matrix. In fact, Perry-Smith [23] argues that creative influence can be classified into two main categories: *strong* and *weak* influence, where *strong* influence, as opposed to *weak* influence, is the result of many redundant connections between two nodes. In other words, the influence of an artist on another artist is stronger if the amount of connections between the two is high. Taking into account long walks results in many connections that eventually turn a weak influence into a strong one. This can also be seen in the example of Figure 2. The influence of *Pete Townshend* on *Halsey* can be safely

ignored. Taking into account walks longer than degree 2 wrongly detects this relationship as a *strong* one. In Equation 2 longer walks are considered less important. However, this proves not to be enough, as the number of distinct longer walks from two entities mitigates this discount and results in less accurate predictions. Further investigation on other converging series used for decaying weights can result in more accurate performances.

**Table 2**

Result from the experiment described in Section 4. Each value is reported alongside its standard deviation.

$w$	MAP	DCG	MRR
Uniform	$0.23 \pm 0.13$	$0.48 \pm 0.13$	$0.43 \pm 0.24$
Frequency	$0.17 \pm 0.12$	$0.43 \pm 0.12$	$0.28 \pm 0.23$
Inverse Frequency	$0.23 \pm 0.16$	$0.49 \pm 0.16$	$0.43 \pm 0.27$
Learned	<b><math>0.28 \pm 0.13</math></b>	<b><math>0.53 \pm 0.11</math></b>	<b><math>0.50 \pm 0.17</math></b>
DNN	$0.1 \pm 0.17$	$0.3 \pm 0.19$	$0.16 \pm 0.27$

Table 2 reports the results obtained from the experiments of Section 4 with  $D = 2$ . Learning the function  $w$  used in Equation 3 leads to the best results on aggregate with respect to all the metrics taken into account. Surprisingly, using a neural network as illustrated in Equation 4 does not result in a definite gain with respect to the simpler model of Equation 1. The main reason for that is the lack of training data. The network is not able to generalise over the target task and tends to overfit in the training data despite the small number of parameters.

**Table 3**

Statistics on the learned weights from the best model of table 2.

Relationship	$\max(w)$	$\text{mean}(w)$	$\text{median}(w)$	$\min(w)$
has acquaintance	9.456889	9.244201	9.286146	8.995128
has bandmate	9.958392	7.067571	6.286661	6.230026
has friend	15.739528	14.435940	14.669381	13.340489
has mentor	4.089674	3.709073	3.636836	3.510836
has pupil	12.359836	10.432499	11.434330	6.429887
is friend of	4.090974	3.709161	3.636191	3.510761

In Table 3 the weights learned by the best method of Table 2 are described. Judging from the mean and median values, the most important relationships are the *friendship* and *mentorship* relationship. The former aligns with the findings of [3]: the influences of famous artists can be largely associated with the connections they have with other meaningful artists from the same clique. Moreover, it is important to notice how the inverse relationship, *is friend of*, has a lower weight when compared to the *has friend* relationship. In order to understand this phenomenon it is sufficient to contextualise it in the task we have identified. An artist can be influenced by a friend only if the artist itself acknowledges that relationship. If the relationship is asymmetric, i.e. one of the two artists is unaware of it, the influence between the two should be much weaker. An interesting phenomenon happens with the *mentorship* relationship. Artists are much more influenced by their students (*has pupil* relationship) rather than their mentors. This

is a direct result of taking into account relationships that span multiple edges rather than a direct connection. Even though an artist can be influenced by one of its students, we argue that the result of the high weight is explained best by a *co-pupil* relationship. An artist is influenced by another artist when they both share the same mentor. This aligns with the findings of [7], where mentors are seen as a *bridge* between two artists.

## 6. Conclusion

In this work, we present a novel method, described in Section 3, to detect creative influence between artists using techniques based on graph theory and complex network science. Our method takes into account the individuality of a relationship type through the use of an ontology illustrated in Section 3.1. By framing the influence prediction task as a classification task we are able to obtain an interpretable model that performs better than robust baselines. The results described in Section 5 highlight how a straightforward combination of the different graph planes identified by the different relationship types results in accurate results. Moreover, the weights assigned to each relationship type are in line with relevant socio-cognitive and psychological findings, thus additionally validating the results. Our attempt to increase the accuracy of our results through the use of a machine learning approach led to less accurate predictions. Nonetheless, it is difficult to objectively rule out the possibility of combining machine learning techniques with our methods. In future works, we plan on extending the dataset available. The main problem with the experiments relying on the neural network can indeed be partially caused by the limited amount of training data for the model. An approach is to employ data augmentation techniques, in order to exploit the data as much as possible and reduce the chances of overfitting the model. With the availability of additional data, for instance by using relation extraction methods [33, 34, 35], more complex architectures can also be used, such as attention-based models [36]. Finally, we would like to explore clustering methods and detect communities of artists within the  $f$ -communicability matrix obtained from the method of Section 3.2. This would enable the identification of cliques of artists that are socially related and hence provide a tool to better understand the creative process of an artist. Combining the relationship types with additional relevant information, such as the socio-cultural context [20], is also an interesting improvement worth investigating. Similarly, combining our approach with the one identified by Saleh et al. [9], where influence between artists is modelled on the basis of the similarity between their artefacts, is an interesting approach as it could help increase the accuracy while also providing examples of artefacts where such similarities can be identified.

## Acknowledgments

The authors would like to thank Daniele Marini for his exploratory work on the subject. This project has received funding from the FAIR – Future Artificial Intelligence Research foundation as part of the grant agreement MUR n. 341.

## References

- [1] G. Hermeren, *Influence in art and literature*, volume 1445, Princeton University Press, 2015.
- [2] I. H. Hassan, The problem of influence in literary history: Notes towards a definition, *The Journal of Aesthetics and Art Criticism* 14 (1955) 66–76. URL: <http://www.jstor.org/stable/426642>.
- [3] B. Mitali, P. L. Ingram, Fame as an illusion of creativity: Evidence from the pioneers of abstract art, HEC Paris Research Paper No. SPE-2018-1305, Columbia Business School Research Paper (2018).
- [4] G. Malanga, *Uptight: The Velvet Underground Story: The Velvet Underground Story*, Omnibus Press, 2009.
- [5] A. Schütz, Making music together: A study in social relationship, *Social research* (1951) 76–97.
- [6] S. Jänicke, J. Focht, Untangling the social network of musicians., in: DH, 2017.
- [7] D. K. Simonton, Artistic creativity and interpersonal relationships across and within generations., *Journal of personality and social psychology* 46 (1984) 1273.
- [8] K. Abe, B. Saleh, A. M. Elgammal, An early framework for determining artistic influence, in: A. Petrosino, L. Maddalena, P. Pala (Eds.), *New Trends in Image Analysis and Processing - ICIAP 2013 - ICIAP 2013 International Workshops*, Naples, Italy, September 9-13, 2013. Proceedings, volume 8158 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 198–207. URL: [https://doi.org/10.1007/978-3-642-41190-8\\_22](https://doi.org/10.1007/978-3-642-41190-8_22). doi:10.1007/978-3-642-41190-8\_22.
- [9] B. Saleh, K. Abe, R. S. Arora, A. M. Elgammal, Toward automated discovery of artistic influence, *Multim. Tools Appl.* 75 (2016) 3565–3591. URL: <https://doi.org/10.1007/s11042-014-2193-x>. doi:10.1007/s11042-014-2193-x.
- [10] A. M. Elgammal, B. Saleh, Quantifying creativity in art networks, in: H. Toivonen, S. Colton, M. Cook, D. Ventura (Eds.), *Proceedings of the Sixth International Conference on Computational Creativity, ICCO 2015, Park City, Utah, USA, June 29 - July 2, 2015*, computationalcreativity.net, 2015, pp. 39–46. URL: [http://computationalcreativity.net/icco2015/proceedings/2\\_3Elgammal.pdf](http://computationalcreativity.net/icco2015/proceedings/2_3Elgammal.pdf).
- [11] K. O’Toole, E. Horvát, Novelty and cultural evolution in modern popular music, *EPJ Data Sci.* 12 (2023) 3. URL: <https://doi.org/10.1140/epjds/s13688-023-00377-7>. doi:10.1140/epjds/s13688-023-00377-7.
- [12] D. Park, J. Nam, J. Park, Novelty and influence of creative works, and quantifying patterns of advances based on probabilistic references networks, *EPJ Data Sci.* 9 (2020) 2. URL: <https://doi.org/10.1140/epjds/s13688-019-0214-8>. doi:10.1140/epjds/s13688-019-0214-8.
- [13] C. Scott, *Music and Its Secret Influence: Throughout the Ages*, Simon and Schuster, 2013.
- [14] M. C. Pattuelli, C. Weller, G. Szablya, Linked jazz: An exploratory pilot, in: T. Baker, D. I. Hillmann, A. Isaac (Eds.), *Proceedings of the 2011 International Conference on Dublin Core and Metadata Applications, DC 2011, The Hague, The Netherlands, September 21-23, 2011, Dublin Core Metadata Initiative, 2011*, pp. 158–164. URL: <http://dcpapers.dublincore.org/pubs/article/view/3637>.
- [15] S. Pozza, F. Tudisco, On the stability of network indices defined by means of matrix

- functions, *SIAM J. Matrix Anal. Appl.* 39 (2018) 1521–1546. URL: <https://doi.org/10.1137/17M1133920>. doi:10.1137/17M1133920.
- [16] C. H. Smith, P. Georges, Composer similarities through “the classical music navigator”: Similarity inference from composer influences, *Empirical Studies of the Arts* 32 (2014) 205–229.
- [17] P. Georges, A. Seckin, Music information visualization and classical composers discovery: an application of network graphs, multidimensional scaling, and support vector machines, *Scientometrics* 127 (2022) 2277–2311. URL: <https://doi.org/10.1007/s11192-022-04331-8>. doi:10.1007/s11192-022-04331-8.
- [18] A. E. White, J. C. Kaufman, M. Riggs, How “outsider” do we like our art?: Influence of artist background on perceptions of warmth, creativity, and likeability., *Psychology of Aesthetics, Creativity, and the Arts* 8 (2014) 144.
- [19] E. Stamkou, G. A. van Kleef, A. C. Homan, The art of influence: When and why deviant artists gain impact., *Journal of Personality and Social Psychology* 115 (2018) 276.
- [20] M. Mauskopf, E. Quintane, N. Askin, J. M. Mol, Embeddedness and the production of novelty in music: A multi-dimensional perspective, *Academy of Management Proceedings* 2017 (2017) 16678. URL: <https://doi.org/10.5465/AMBPP.2017.2>. doi:10.5465/AMBPP.2017.2. arXiv:<https://doi.org/10.5465/AMBPP.2017.2>.
- [21] K. J. Borowiecki, Good reverberations? teacher influence in music composition since 1450, *Journal of Political Economy* 130 (2022) 991–1090.
- [22] J. E. Perry-Smith, C. E. Shalley, The social side of creativity: A static and dynamic social network perspective, *Academy of management review* 28 (2003) 89–106.
- [23] J. E. Perry-Smith, Social yet creative: The role of social relationships in facilitating individual creativity, *Academy of Management journal* 49 (2006) 85–101.
- [24] M. C. Pattuelli, K. Hwang, M. Miller, Accidental discovery, intentional inquiry: Leveraging linked data to uncover the women of jazz, *Digit. Scholarsh. Humanit.* 32 (2017) 918–924. URL: <https://doi.org/10.1093/llc/fqw047>. doi:10.1093/llc/fqw047.
- [25] J. Lethem, *The ecstasy of influence* (2007).
- [26] S. Borgo, R. Ferrario, A. Gangemi, N. Guarino, C. Masolo, D. Porello, E. M. Sanfilippo, L. Vieu, DOLCE: A descriptive ontology for linguistic and cognitive engineering, *Appl. Ontology* 17 (2022) 45–69. URL: <https://doi.org/10.3233/AO-210259>. doi:10.3233/AO-210259.
- [27] A. Krisnadhi, F. Maier, P. Hitzler, OWL and rules, in: A. Polleres, C. d’Amato, M. Arenas, S. Handschuh, P. Kroner, S. Ossowski, P. F. Patel-Schneider (Eds.), *Reasoning Web. Semantic Technologies for the Web of Data - 7th International Summer School 2011, Galway, Ireland, August 23-27, 2011, Tutorial Lectures*, volume 6848 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 382–415. URL: [https://doi.org/10.1007/978-3-642-23032-5\\_7](https://doi.org/10.1007/978-3-642-23032-5_7). doi:10.1007/978-3-642-23032-5\_7.
- [28] J. A. Bondy, U. S. R. Murty, *Graph Theory*, Graduate Texts in Mathematics, Springer, 2008. URL: <https://doi.org/10.1007/978-1-84628-970-5>. doi:10.1007/978-1-84628-970-5.
- [29] L. Katz, A new status index derived from sociometric analysis, *Psychometrika* 18 (1953) 39–43.
- [30] E. Estrada, N. Hatano, Communicability graph and community structures in complex networks, *Appl. Math. Comput.* 214 (2009) 500–511. URL: <https://doi.org/10.1016/j.amc>.

2009.04.024. doi:10.1016/j.amc.2009.04.024.

- [31] S. Bruch, An alternative cross entropy loss for learning-to-rank, in: J. Leskovec, M. Grobelnik, M. Najork, J. Tang, L. Zia (Eds.), WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021, ACM / IW3C2, 2021, pp. 118–126. URL: <https://doi.org/10.1145/3442381.3449794>. doi:10.1145/3442381.3449794.
- [32] F. Orabona, T. Tommasi, Training deep networks without learning rates through coin betting, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 30, Curran Associates, Inc., 2017. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/7c82fab8c8f89124e2ce92984e04fb40-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/7c82fab8c8f89124e2ce92984e04fb40-Paper.pdf).
- [33] P. H. Cabot, R. Navigli, REBEL: relation extraction by end-to-end language generation, in: M. Moens, X. Huang, L. Specia, S. W. Yih (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021, Association for Computational Linguistics, 2021, pp. 2370–2381. URL: <https://doi.org/10.18653/v1/2021.findings-emnlp.204>. doi:10.18653/v1/2021.findings-emnlp.204.
- [34] Y. Lu, Q. Liu, D. Dai, X. Xiao, H. Lin, X. Han, L. Sun, H. Wu, Unified structure generation for universal information extraction, in: S. Muresan, P. Nakov, A. Villavicencio (Eds.), Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022, Association for Computational Linguistics, 2022, pp. 5755–5772. URL: <https://doi.org/10.18653/v1/2022.acl-long.395>. doi:10.18653/v1/2022.acl-long.395.
- [35] S. Wadhwa, S. Amir, B. C. Wallace, Revisiting relation extraction in the era of large language models, in: A. Rogers, J. L. Boyd-Graber, N. Okazaki (Eds.), Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023, Association for Computational Linguistics, 2023, pp. 15566–15589. URL: <https://doi.org/10.18653/v1/2023.acl-long.868>. doi:10.18653/v1/2023.acl-long.868.
- [36] A. Galassi, M. Lippi, P. Torrioni, Attention in natural language processing, IEEE Trans. Neural Networks Learn. Syst. 32 (2021) 4291–4308. URL: <https://doi.org/10.1109/TNNLS.2020.3019893>. doi:10.1109/TNNLS.2020.3019893.