

Modelling Interestingness: a Workflow for Surprisal-based Knowledge Mining in Narrative Semantic Networks

Cosimo Palma^{1,2}

¹University of Naples "L'Orientale", Via Duomo, 219, 80139 Napoli, Italy

²University of Pisa, Lungarno Antonio Pacinotti, 43, 56126 Pisa, Italy

Abstract

This working paper outlines ongoing and planned efforts aimed at achieving an objective modelling of *interestingness* in cross-domain knowledge bases. In pursuit of this objective, clickstream data serves as a primary component for developing a novel measure of entity-related popularity. This measure is then integrated with two couple-related similarity measures, culminating in the formulation of a new interestingness law. This principled formalization is designed to undergo human validation, ultimately enhancing its reliability and comprehensiveness. The present contribution is intended to be propaedeutic to the development of a pipeline having a Knowledge Graph as input, and an expanded version of the same as output, whereby every link is labelled by an interestingness score, thus highlighting the most interesting paths, determined according to the proposed domain-specific heuristics for interestingness detection. This work is expected to yield significant benefits for Automatic Story Generation. Although this discipline, aided by Machine Learning, has made remarkable progress in surface-level text realization, it still grapples with producing qualitatively rich outputs that offer substantive informativeness. To address this challenge, a Knowledge Graph (particularly its most compelling paths identified through the proposed methodology) is anticipated to integrate the Large Language Model, thus harnessing the final output with the contextual information selected by users throughout the entire workflow- a scenario which is particularly valuable in educational settings, where generated stories frequently serve pedagogical purposes.

Keywords

Interestingness measures, Computational Creativity, Knowledge Mining, Automatic Story Generation

1. Delineation of knowledge gap and problem statement

Knowledge graphs are notably the most used and intuitive data structure for modelling narratives [1, 2, 3]. For what concerns Story Generation, are indeed the LLMs representing nowadays the most valuable resource [4, 5]. Non trivial issues, such as the enhancement of *coherence* [6], *plausibility* [7] and *contextuality* [8] have already been successfully addressed, not rarely leveraging KGs. Beside *coherence*, *interestingness* represents one of the most common metrics for evaluation of synthetic stories [9].¹ According to the way they are currently modelled, there cannot be but a trade-off between the two: the coherence of the represented events hinders the interestingness, which is conversely boosted by unusual yet still sound connections between textual units, whether they are events, sentences or terms.

Beyond paving the way to the solution of this problem, the real-life scenario from which the semi-automatic pipeline mentioned in this paper has drawn its motivation is the work usually performed by essayists and novelists who are used to exploit archives, libraries, or any other type of cultural repository as a source for inspiration and material.² Nowadays such collections are often digitized and available to the public as Linked Open Data (LOD), and therefore displayable and processable as KGs.

SEMMES'24: Semantic Methods for Events and Stories co-located with 21th Extended Semantic Web Conference (ESWC2024), May 27, 2024, Hersonissos, Greece.

✉ cosimo.palma@phd.unipi.it (C. Palma)

🌐 <https://glottocrisio.github.io/bupbengoc/> (C. Palma)

🆔 0000-0002-8161-9782 (C. Palma)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹The metrics listed therein beside coherence are: relevance, empathy, surprise and engagement. Although interestingness is not explicitly mentioned, the latter three are considered as interestingness-based metrics.

²The retrieved information would then represent what afterwards is "shown not told", as the adagio of Fiction Writing recites. The creation of a plot, which naturally reuses and lets the so-retrieved information *emerge*, will pertain to a following step of implementation.

The process of digitization is usually very costly in terms of time and resources, and mainly undertaken for *preservation* reasons [10]. However, *reuse* from society and the economic gain attached to it, is what would really catalyse this effort. In turn, society would be highly encouraged to exploit such resources if they are abundant, open, clean and duly maintained. It follows that a highly beneficial aim is the creation of tools aiming at fuelling and smoothing out this virtuous cycle involving resources and society. For this reason, the exploitation of the large availability of LOD to build the Knowledge base required in our case study is regarded as particularly significant.

Shifting the interestingness modelling from the text-surface realization of synthetic stories towards the general story plotting, is individuated as a straightforward way to detangle oneself from the initial apparent contradiction between coherence and interestingness. For this reason, above all, the detection of *interesting* paths in cross-domain *Knowledge Graphs* (KG) has to be considered a task of absolute importance in the domain of Computational Creativity, particularly in Automatic Story Generation (ASG). The almost ubiquitous definition of "interestingness" in current literature equates to that of "relevance", thus relating to probabilistic measures or other heuristics intended to capture co-occurrence and *similarity* between two or more concepts [11]. Conversely, this contribution intends to address the problem of detecting interesting *paths*³ in a Knowledge Graph tailored for narrative purposes (such as, for instance, Event KGs [12, 3]).

To this extent, a novel formulation of *interestingness* is required. The present modelling work underlies the implementation of an application for further expanding Knowledge Graphs (see Fig. 1) retrieved by the HILDEGARD workflow [13] (i.e., semi-automatically elaborated on two or more entry-seeds constituted by heritage objects as retrieved in Digital Heritage databases) and extracting the most interesting paths.

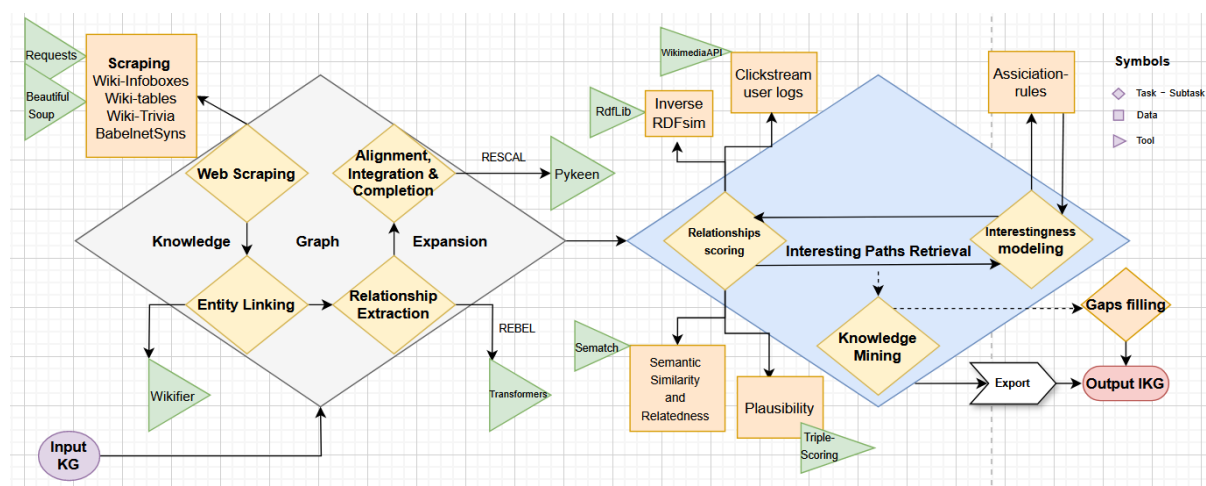


Figure 1: ISAAKX - Interestingness-grounded Semi-AutomAtic Knowledge eXpansion. A pipeline proposal for an Interestingness-based Knowledge Graph, with related implementation tools and methods.

In turn, this module is intended as the second one of a broader pipeline, featuring in its third step a module which automatically generates from the retrieved RDF-triples SVO (Subject-Verb-Object) triples, that will in turn be prompted in the final language model according to a Chain-of-Thoughts [14] paradigm to perform an *informed* ASG.

2. A brief literature review on Interestingness

Humans present an inherent need for novelty probably associated with the dopamine D4 receptor gene [15]. Recent findings in psychology confirm that surprise is summoned by unexpected (schema-discrepant) events and its intensity is determined by the degree of schema-discrepancy [16]. Intuitively,

³In graph theory, a path is a walk in which neither edges nor nodes are repeated. A path ending with the same node it started, is called "cycle".

humans find interesting what is not obvious, yet at the same time not random, since randomness might cause confusion and boredom.

In both education and narrative, as well as in Information retrieval, the *analogy* represents an outstanding device to create interesting associations: the issue of overlapping concepts throughout different media and the human senses reserved for their fruition, is addressed from a cognitive as well as computational perspective in the *Conceptual Blending* theory [17]. According to it, a process starts by finding a partial mapping between elements of two input spaces that are perceived as analogous with respect to their graph representation. Afterwards the so-called generic mental space encapsulates the conceptual structure shared by the input spaces, generalising and possibly enriching them. This space provides guidance to the next step of the process, where elements from each of the input spaces are selectively projected into a new mental space, called the *blend space*.

Other devices generally regarded as interesting are, for instance, jokes, aphorisms, trivia, poetries, novels and so on. The reason why we find them interesting, can be boiled down to the concept of *twist*. The twist individuated in the following is the gap between two different kinds of similarity, but plenty of other ones can be eventually found, such as *paradoxes* and *contradictions*, or even graph- structural gaps [18]. Before proceeding with the proposed method to pinpoint the most "interesting" triples in a KG, let us first review the existing scientific literature on the topic of interestingness (also known as novelty/surprisingness) measures.

The concept of *interestingness* has been differently defined in accordance to the discipline it occurs in. In Hilderman & Hamilton (1999) [19] a thorough survey of all possible measures is performed in the field of Knowledge Discovery (KD), where they are ranked by *representation* (the dataset format on which they are to be applied, such as classification rules, summaries, or association rules), *foundation* (probabilistic, distance-based, syntactic or utilitarian), *scope* (single rule/rule set) and *class* (objective/subjective).

Subjective measures involve the user's background knowledge about the data, encompassing input bias, constraints, beliefs, expectations, or interactive feedback. These subjective measures are typically integrated into the mining process⁴, and their representation can vary. In contrast, objective interestingness measures solely rely on the data itself, without requiring additional user inputs. Our approach, specified in the section "Task's formal definition", aims at *objectivising* the subjective measures such as "unexpectedness" and "novelty" [20]. It could be argued that subjective measures of interestingness recognize that a pattern may be interesting for someone, and for somebody else not. For example, a pattern discovering some security trading irregularities, such as insider trading, may be of great interest to the officials from the Securities and Exchange Commission and of very little use to a homeless person living in Naples [21]. The subjectivity in this paper it is advocated for is rather "Intersubjectivity", i.e. objective subjectivity, where with "objective" is typically meant *conventional*. One of the most relevant measures to this respect is the Silbershatz and Tuzhilin's Interestingness, which determines the extent to which a *soft belief* is changed as a result of encountering new evidence⁵

Large foundational Upper Ontologies, such as *DOLCE* and *Cyc* can provide the encoding of general knowledge, the *hard beliefs*. Also domain-specific Ontologies, such as CIDOC-CRM⁶ for Cultural Heritage, often contain in their schema-specification formulas in Description Logic, describing inferences can be used to further expand, in a downstreaming fashion, the hard beliefs required to proceed in modelling interestingness. As an example for the so-called "objective" measures, may it suffice to report Freitas' Surprisingness [23], according to which the interestingness of discovered knowledge is obtained via the explicit detection of occurrences of Simpson's paradox.⁷ According to our approach,

⁴Knowledge discovery in databases, commonly also referred to as knowledge mining, involves the effective identification of previously undiscovered, valid and potentially valuable patterns within extensive databases. It encompasses diverse techniques and algorithms, each varying in the types of data analysed and the method of representing the acquired knowledge.

⁵A *soft belief* is one that an agent is can easily change provided that new evidence is encountered [19]. In [22] subjective measures of interestingness are considered in more depth. These measures are classified into *actionable* and *unexpected*, and the relationship between them is examined.

⁶<https://cidoc-crm.org/Version/version-7.1.3>.

⁷Simpson's Paradox occurs when a trend that is visible within separate groups of data vanishes or reverses once the groups are merged. This phenomenon happens due to the presence of a lurking variable or confounding factor that affects the relationship between the observed variables. As a result, combining the data can lead to misleading or counterintuitive

the *paradox/contradiction* is identified as a heuristics that effectively capture interestingness, which similarly to the analogy links entities very similar to some extent, but very different to another one. The allocated Interestingness measures are meaningful in a work frame where *transactions* happens, i.e. operations that modify the database. Transactional databases are used, for instance, in banking or online retail systems. Association rules are patterns or relationships discovered within transactional databases that reveal connections between different items based on their co-occurrence in transactions. These changes can form *patterns*, that association rules are then supposed to capture, allowing the detection of some interesting yet hidden ones. These rules are often represented in the form of "if-then" statements. For example, a simple association rule might be: "If a customer buys item A, then he is likely to buy item B as well."

Rules interestingness involves assessing the value or relevance of these discovered rules, usually against ground-truth baselines. Various measures are used to determine the interestingness of association rules, including (freely adapted from [11]):

1. **Support:** The frequency that A and B co-occur in a single transaction;
2. **Confidence:** The percentage of transactions containing both A and B compared to the number of transactions containing only B;
3. **Lift:** The ratio of the probability of A and B co-occurring in the same transaction compared to the probability expected if A and B were independent;

Although the attempt presented in this contribution proposes not a "more interesting" association rule, rather an rule for interestingness, these association rules could be indeed exploited to assess the relationships among the individuated sub-components, thus helping in polishing the general interestingness law (2). However, the same and other papers reporting similar measures, apply them in relational databases without disruptions, interpreting the above-mentioned As and Bs not as items or item-sets, but rather as *entities*. If the similarity between temporal Knowledge Graphs and transactional databases is intuitive, it is required, regarding static Knowledge Graphs, to introduce other concepts, such as, for example, the *subpatterns* and *superpatterns* mentioned in [11]. Therein, *Surprisingness I* evaluates how difficultly a pattern's frequency can be derived from its components (subpatterns); *Surprisingness II* evaluates how difficultly a pattern's frequency can be derived from its superpatterns. The authors define *potentially interesting patterns* those ones showing high *Surprisingness I*, and just high enough occurrence to catch attention (but not as high as *well-known patterns*, nor as low as *exceptions*). They result potentially interesting because they are not very well-known yet [ibid.].

Despite this approach, whose mindset constitutes considerable part of the inspiration of the present work, all considered literature on interestingness measures adopted on Networks, including the one referring explicitly to ontology based data [24, 22], do not take in consideration that the relationships may also be weighted, but use the concept of Semantic Similarity, realized in *Rada* similarity, exploiting the shortest path between two entities, and *Resnik* similarity, where the similarity between classes x and y is defined as the information content of their Most Informative Common Ancestor [ibid.]. Even in that case, interestingness is tuned on similarity. Although quite often the Knowledge based used for analysis and evaluation includes large Cross-domain Knowledge Graphs, such as *Wikidata*, to the best of our knowledge the *clickstream* data feature is never leveraged to attempt at a more comprehensive measurement of interestingness, which in this paper is regarded as an objective measure based on *novelty*⁸.

In [25] exploiting user browsing behavior in clicks from one page to another has been already successfully implemented in the domain of Natural Language Processing (NLP), but without taking into account the complexity of general interestingness.

outcomes, where the overall association may be opposite to that observed within individual subgroups.

⁸A glimpse on recent use of clickstream-data can be found at <https://wikimediafoundation.org/news/2018/01/16/wikipedia-rabbit-hole-clickstream/>.

3. Task's formal definition

The modelling of "Interestingness" depends, among other factors, on the magnitude of application. For the moment, we consider only the assessment of the interestingness of a relationship between two entities. An entity belonging to a Cross-domain Knowledge Graph, such as *Wikipedia*, can be classified according to measures that are easier to model mathematically, such as, for instance, the *popularity*. Intuitively, the more accesses a web pages receives, the more popular it is. Nevertheless, if two entities have the same number of views, the one with fewer incoming links is deemed more popular because accessing it is less straightforward. To capture this heuristic we resort to the concept of *graph centrality*. The literature presents many measure for assessing node centrality in a given graph. *Subgraph centrality* of a node u in G can be found using a spectral decomposition of the adjacency matrix [26]: $C_{\text{subgraph}(i)} = \sum_{j=1}^N (v_j^i)^2 e^{\lambda_j}$ where v_j is an eigenvector of the adjacency matrix A of G corresponding to the eigenvalue λ_j . Subgraph centrality provides a measure of how well-connected or influential a specific set of nodes is within the overall network structure. This centrality measure has been preferred to others because it enables the assessment of influence within particular substructures of a network, offering a more focused perspective on centrality. This approach aids in the recognition of nodes that, while lacking high global centrality, play a critical role within specific local communities or subgraphs. For the moment, let us rely just on popularity, as absolute measure with whom nodes can be labelled.

This example of a basic heuristic for the popularity \mathfrak{P} of an entity i , is rendered as in the following:

$$\mathfrak{P}(i) \simeq \frac{\text{clickstream}(i)}{C_{\text{subgraph}(i)}} \quad (1)$$

The *clickstream* data on single Wikipedia-entities can be collected by means of Python libraries such as MediaWikiAPI.⁹ On the other hand, to fetch the clickstream-data related to entity couples, there is no other way than directly exploiting Wikimedia Clickstream Data Dumps.¹⁰, as performed in the project *WikiNav*.¹¹ Although the sketched formalization gives priority to the clickstream data, which also represents the main resource proposed in this contribution, its scarcity raises some concerns: *Wikipedia* seems to be the only Knowledge Graph presenting its availability. Conversely, graph centrality is a structural property of every graph. For this reason, exploring the correlation between clickstream-data and graph-centrality might lead to a satisfactory approximation of popularity leveraging only centrality, thus allowing its detection on every graph, without requiring clickstream-data. For the moment, since the selected graph does present such a resource, graph centrality is here used as a coefficient, to adjust popularity according to the formulated heuristics.

Among *relative measures*, *similarity* is definitely the most known in literature. Two types of similarity between two entities have been identified: *corpus*- and *knowledge*-based similarity. According to our heuristics for interestingness as contraddiction, that two entities can be similar in one way, yet dissimilar in the other, would indeed result interesting. Capturing the precise gap of this interestingness would deliver its precise degree of interestingness. Both measures can be extracted by means of the Python package *Sematch* [27].¹² Knowledge-based similarity is therein captured by their Relatedness, whose mathematical formalization has been expressed with:

$$\text{relatedness}(a, b) = \frac{\log(\max(|A|, |B|)) - \log(|A \cap B|)}{\log(|W|) - \log(\min(|A|, |B|))}$$

,where a and b are the two articles of interest, A and B are the sets of all articles that link to a and b respectively, and W is set of all articles in Wikipedia. The relatedness of a candidate sense is the weighted average of its relatedness to each context article, where the weight of each comparison is defined in the next section [28]. Otherwise, it would also be possible by means of *Rdf-similarity* [29]. Furthermore, corpus similarity is calculated in *Sematch* extracting YAGO concepts of the entity from

⁹<https://pypi.org/project/mediawikiapi/>.

¹⁰<https://dumps.wikimedia.org/other/clickstream/readme.html>.

¹¹<https://wikinav.toolforge.org/>.

¹²<https://pypi.org/project/sematch/>.

DBpedia using *EntityFeatures* class. Then, the top five concepts with highest graph-based information contents are selected and composed as concept list. Finally, corpus similarity of two entities is calculated by Pointwise Mutual Information or Latent Semantic Analysis of the related concept word lists, based on large corpora [30]. Although YAGO concepts can equate to terms, they are not much likely to occur in large corpora as they are. An effective work-around to this drawback would be using the more general *enLabel* class obtained by using the *Wikifier*. Alternatively, pretrained word embeddings can be exploited to vectorized the obtained word lists, and a cosine similarity can be then performed between the two vectors. An other option to calculate the similarity of two vectors is executing their dot product, and divide it by the product of their (for instance, Euclidean) norm.

It can be argued that concepts can be often be treated as entities, which cannot be said of the converse. In case also a conceptual similarity between two specific entities is needed, it is possible to extract concepts strictly related to entities by leveraging the *enLabel* property of the *Wikifier*¹³, the Entity Linking tool used also in the previous step of the pipeline. After extracting the word-vectors of concepts related to the entity, the *average vector* for each entity would represent a shallow representation of the concept in the vector space. They can be then compared, by means of similarity measures such as the Cosine Similarity. The obtained value would be a meaningful representation of the similarity between the original entities.

Now that the general concept of *interestingness* has been broken down in other components which can be more easily modelled mathematically, let us try to leverage them to possibly move towards a modelling of interestingness. Since an entity/node can be either popular or unpopular and a relationship as corpus- or knowledge based, our problem can be modelled as a permutation with repetition: if the node can be of two types and the relationship of eight types (namely, high corpus- and high knowledge-based similarity, low corpus- and low knowledge-based similarity, high corpus - and low knowledge-based similarity, low corpus- and high knowledge-based), we can have: $2 \times 8 \times 2 = 32$ possibilities, among which we have selected the following ones as interesting, because the only ones showing the paradox/contradiction characteristic mentioned at the beginning:

1. Popular Entity (-) HIGH corpus- AND HIGH knowledge-based similarity (-) Unpopular Entity;
2. Popular Entity (-) HIGH corpus- BUT LOW knowledge-based similarity (-) Unpopular Entity;
3. Popular Entity (-) HIGH knowledge- BUT LOW corpus-based similarity (-) Unpopular Entity (e.g. Trivia);
4. Popular Entity (-) HIGH corpus- BUT LOW knowledge-based similarity (-) Popular Entity
5. Popular Entity (-) HIGH knowledge- BUT LOW corpus-based similarity (-) Popular Entity

Given all aforementioned information, an attempt of interestingness model can be individuated in the following formula:

$$\mathcal{I}_{n1,n2} \cong \frac{\mathfrak{P}_{n1,n2} \times |\mathfrak{S}_{n1,n2} - \text{DBpediaRel}_{n1,n2}|}{\log_{10}(\text{clickstream}_{n1,n2})} \quad (2)$$

where:

$$\mathfrak{S}_{n1,n2} \cong \ln \left(\left| \frac{(\text{CosineSim}_{n1,n2} + \text{DBpediaSim}_{n1,n2})}{2} \right| \right) \text{ and: } \mathfrak{P}_{n1,n2} \cong \ln (\mathfrak{P}_{n1+n2} + |\mathfrak{P}_{n1-n2}|)$$

The distributional similarity between two entities is obtained through the simple average between cosine similarity (corpus-based) and DBpedia similarity (as in *Sematch*, concepts- or labels- based), whereas the popularity has been modelled in order to capture the following set of constraints:

1. The overall interestingness increases if both entities have a high popularity;
2. The overall interestingness increases if one entity is considerably more popular than the other;

The clickstream has been set as the denominator of the main law (2) according to the interpretation that a lower clickstream, hence a lesser obviousness of the link, shall produce a higher overall interestingness

¹³<https://wikifier.org/>.

value. Once the general interestingness of every triple is computed, one can proceed on the second step of our quest: the search of most interesting paths. As explained in [31], on a diachronic dimension, interestingness cannot be simply formalized as a series of interesting events. Therefore, our problem cannot be simply formalized as a maximization problem, of the kind: find the path of n given directly-linked elements, whose relationships score is the highest. In this case we need a broader and at the same time deeper consideration of what is to be considered interesting, on a sequential level. Following the *magic squares model* of interestingness [ibid.], our problem can be formalized, for instance, in a similar fashion:

Find the sequence of a given number n of relationships, whose interestingness scores, once arranged in a squared matrix, most resemble the properties and constraints of a magic square of order \sqrt{n} .

ClickstreamE1E2	PopularityE1	PopularityE2	PopularityDiff	PopularitySum	CosineSimilarityE1E2	DBpediaSimilarityE1E2	DBpediaRelatednessE1E2	InterestingnessE1E2
15	1831	113391.33	111560.33	115222.33	0.26	0.08	0.5	14.3
23	80818.67	113391.33	32572.66	194210	0.39	0.14	0.53	8.48
591	14588.98	113391.33	98802.35	127980.31	0.59	0.33	0.26	1.18
15	30504.74	113391.33	82886.59	143896.07	0.37	0.09	0.44	9.67
11	117113.39	113391.33	3722.06	230504.72	0.44	0.09	0.3	5.28
32	2330	113391.33	111061.33	115721.33	0.32	0.33	0.37	3.27
14	2592.67	113391.33	110798.66	115984	0.32	0.12	0.41	9.78
40	15062.33	113391.33	98329	128453.66	0.56	0.09	0.3	1.73
35	17000	113391.33	96391.33	130391.33	0.39	0.1	0.37	5.62
22	16401	113391.33	96990.33	129792.33	0.45	0.14	0.41	5.44
25	113472.17	113391.33	80.84	226863.5	0.28	0.14	0.56	10.92
30	46281.9	113391.33	67109.43	159673.23	0.39	0.12	0.46	7.1

Table 1

A glimpse of the features extracted calculated on a sample entity-couples dataset according to the proposed modelling.

4. Discussion of results and future work

Throughout this paper a modelling of "Interestingness" towards "Novelty" (instead of the usual "Relevance") has been fleshed out, addressing the theoretical foundation of the task "Interesting paths retrieval", and its sub-tasks (see Fig. 1). The dataset to perform the required analysis has been harvested by means of the off-the-shelf tools (see table 1)¹⁴. This step, achieved as preliminary proof of concept, has failed to show unequivocal correlations among features at a first overview. Before proceeding with a deeper analysis, it is first of all needed to validate the obtained values, as well as to clean and properly normalize them. Furthermore, the last word on this topic cannot be said without involving human evaluation, by designing reliable score-sheets and tests that can steer the formalization towards a standardized and shared effort. To this extent, annotation tools such as the *Amazon Mechanical Turk* can be used to ask more workers to perform the task: "Label this entity-couple as interesting or not interesting, whereby "interesting" means that you would like to understand how precisely these entities are related". The confidence of the annotation would then represent the real value of interestingness, comprised between 0 and 1. After collecting the ground-truth values, it will be then necessary to discover patterns among the features and their correlations with the ground-truth values.

At this point it can be argued, from a point of view of mere practicality, that if the problem can be solved in a neural fashion, the interestingness law is not required any more. However, in the domain of Computational Creativity, the understanding of mental processes under both neural and symbolic paradigmas, treating data as rules and rules as data, would allow their integration and fibring, opening new avenues for the implementation of *truly* intelligent systems, more effectively mimicking human cognitive processes. As it can be noticed, in our proof of concept only couples of entities have been taken in consideration, without their relationship. In our case, intuitively, the entities must be linked by the Dbpedia Ontology relationship "dbo:wikiPageWikiLink". Specifying more precise relationships by means of extraction methods, may represent an alternative avenue for narrativising knowledge graphs (see Fig. 1).

¹⁴The code related to the project *WikiWooW* can be freely accessed at <https://github.com/Glottocrisio/WikiWooW>.

The first part of the pipeline, as to be observed on the left side of the same figure, requires the KG to be expanded by integration with general KGs in a semi-automatic fashion. The interestingness-values for every link will then be assigned according to the same function previously identified for triple scoring. To adapt the system to the subjectivity of the user, a gradient bar can be offered as input for every parameter, and according to his own grading of each one the parameters, it will produce a different rule, or a differently weighted rule.

Beside Wikipedia Infoboxes and Trivia, the starting knowledge base can be furtherly expanded already in an "interesting" way: rule-based approach can be valuable for ambiguity detection, as also for *blends* retrieval, according to the aforementioned conceptual blending theory [17]. More synthetically, an alignment with Framester [32], a large multimodal knowledge base including even Sentiment-based graphs, is supposed to be beneficial for our research direction. The eventual finding of a consistent and shared formalization of "interesting path", as hinted in the previous paragraph, can lead to a stable algorithm for interestingness-based walks, which can be leveraged for graph-embeddings.

References

- [1] I. Blin, Building Narrative Structures from Knowledge Graphs, in: P. Groth, A. Rula, J. Schneider, I. Tiddi, E. Simperl, P. Alexopoulos, R. Hoekstra, M. Alam, A. Dimou, M. Tamper (Eds.), *The Semantic Web: ESWC 2022 Satellite Events*, volume 13384, Springer International Publishing, Cham, 2022, pp. 234–251.
- [2] R. Porzel, M. Pomarlan, L. Spillner, J. Bateman, T. Mildner, C. Santagiustina, Narrativizing Knowledge Graphs, in: *Proceedings of the International Workshop on Knowledge Graph Summarization*, 2022.
- [3] S. Gottschalk, E. Demidova, EventKG - the Hub of Event Knowledge on the Web - and Biographical Timeline Generation, 2019.
- [4] Z. Lin, M. Riedl, Plug-and-Blend: A Framework for Controllable Story Generation with Blended Control Codes, in: *Proceedings of the Third Workshop on Narrative Understanding*, Association for Computational Linguistics, Virtual, 2021, pp. 62–71.
- [5] T. Pradyumna, D. Murtaza, M. Animesh, M. Lara J., H. Brent, O. R. Mark, Controllable neural story generation via reinforcement learning, 2018.
- [6] A. Fan, M. Lewis, Y. Dauphin, Hierarchical Neural Story Generation, in: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, 2018.
- [7] X. Peng, K. Xie, A. Alabdulkarim, H. Kayam, S. Dani, M. Riedl, Guiding neural story generation with reader models, in: Y. Goldberg, Z. Kozareva, Y. Zhang (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2022*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 7087–7111.
- [8] X. Yang, I. Tiddi, Creative Storytelling with Language Models and Knowledge Graphs, in: *Proceedings of the CIKM 2020 Workshops*, Galway, Ireland, 2020, p. 9.
- [9] C. Chhun, P. Colombo, F. M. Suchanek, C. Clavel, Of human criteria and automatic metrics: A benchmark of the evaluation of story generation, in: N. Calzolari, C.-R. Huang, H. Kim, J. Pustejovsky, L. Wanner, K.-S. Choi, P.-M. Ryu, H.-H. Chen, L. Donatelli, H. Ji, S. Kurohashi, P. Paggio, N. Xue, S. Kim, Y. Hahm, Z. He, T. K. Lee, E. Santus, F. Bond, S.-H. Na (Eds.), *Proceedings of the 29th International Conference on Computational Linguistics*, International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 2022, pp. 5794–5836.
- [10] A. Backman, M. Smith, SUSTAINABILITY AND BEST PRACTICES FOR LINKED DATA HERITAGE RESOURCES: SOME CASE STUDIES FROM SWEDEN, Arc Humanities Press, 2023, pp. 159–184.
- [11] S. Ke, P. Spronck, B. Goertzel, A. Van der Peet, Surprisingness - a novel objective interestingness measure in hypergraph pattern mining from knowledge graphs for common sense learning, in: *2021 IEEE International Conference on Big Knowledge (ICBK)*, 2021, pp. 54–61.

- [12] E. Boschee, J. Lautenschlager, S. O'Brien, S. Shellman, J. Starz, M. Ward, ICEWS Coded Event Data, 2015.
- [13] C. Palma, HILDEGARD - Human-In-the-Loop Data Extraction and Graphically Augmented Relation Discovery, *Journal of Computational and Cognitive Engineering* (2024). Forthcoming. Now available at https://www.researchgate.net/publication/378463201_HILDEGARD_Human-In-the-Loop_Data_Extraction_and_Graphically_Augmented_Relation_Discovery.
- [14] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, D. Zhou, Chain-of-thought prompting elicits reasoning in large language models, in: *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Curran Associates Inc., Red Hook, NY, USA, 2024.
- [15] L. Itti, P. Baldi, Bayesian surprise attracts human attention, in: *Advances in neural information processing systems*, 2006, pp. 547–554.
- [16] R. Reisenzein, G. Horstmann, A. Schützwohl, The Cognitive-Evolutionary Model of Surprise: A Review of the Evidence, *Topics in Cognitive Science* 11 (2019) 50–74.
- [17] M. Turner, G. Fauconnier, *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*, Basic Books; Reprint Edition, 2003.
- [18] S. Mellace, K. Vani, A. Antonucci, Relation clustering in narrative knowledge graphs (2020).
- [19] R. J. Hilderman, H. J. Hamilton, Knowledge Discovery and Interestingness Measures: A Survey, *Computer Science* (1999) 28.
- [20] K. MCGARRY, MCGARRY, K.: A survey of interestingness measures for knowledge discovery. *know. eng. rev.* 20(01), 39–61, *Knowledge Eng. Review* 20 (2005) 39–61.
- [21] A. Silberschatz, A. Tuzhilin, What makes patterns interesting in knowledge discovery systems, *IEEE Transactions on Knowledge and Data Engineering* 8 (1996) 970–974.
- [22] An interestingness measure for knowledge bases, *Engineering Science and Technology, an International Journal* 43 (2023) 101417. Publisher: Elsevier.
- [23] A. A. Freitas, On objective measures of rule surprisingness, in: J. G. Carbonell, J. Siekmann, G. Goos, J. Hartmanis, J. Van Leeuwen, J. M. Żytkow, M. Quafafou (Eds.), *Principles of Data Mining and Knowledge Discovery*, volume 1510, Springer Berlin Heidelberg, Berlin, Heidelberg, 1998, pp. 1–9. Series Title: *Lecture Notes in Computer Science*.
- [24] A. C.B, K. Mahesh, Ontology-based data interestingness: A state-of-the-art review, *Natural Language Processing Journal* 4 (2023).
- [25] M. Gamon, A. Mukherjee, P. Pantel, Predicting interesting things in text (2014). URL: <https://api.semanticscholar.org/CorpusID:14414145>.
- [26] E. Estrada, J. A. Rodríguez-Velázquez, Subgraph centrality in complex networks, *Physical Review E* 71 (2005).
- [27] G. Cheng, K. Gunaratna, A. Thalhammer, H. Paulheim, M. Voigt, R. García, Sematch: Semantic entity search from knowledge graph, in: G. Cheng, K. Gunaratna, A. Thalhammer, H. Paulheim, M. Voigt, R. García (Eds.), *Joint Proceedings of the 1st International Workshop on Summarizing and Presenting Entities and Ontologies and the 3rd International Workshop on Human Semantic Web Interfaces (SumPre 2015, HSWI 2015)*, Portoroz, Slovenia, 2015.
- [28] D. Milne, I. H. Witten, Learning to link with Wikipedia, in: *Proceedings of the 17th ACM Conference on Information and Knowledge Management, ACM*, 2008.
- [29] M. Chatzakis, M. Mountantonakis, Y. Tzitzikas, Rdfsim: Similarity-based browsing over dbpedia using embeddings, *Information* 12 (2021) 440.
- [30] R. Mihalcea, C. Corley, C. Strapparava, Corpus-based and knowledge-based measures of text semantic similarity, in: *AAAI*, volume 6, 2006.
- [31] C. Palma, Modelling interestingness: Stories as L-Systems and Magic Squares, in: *Text2Story@ECIR*, 2023. URL: <https://api.semanticscholar.org/CorpusID:258334754>.
- [32] A. Gangemi, M. Alam, L. Asprino, V. Presutti, D. R. Recupero, Framester: A Wide Coverage Linguistic Linked Data Hub, in: E. Blomqvist, P. Ciancarini, F. Poggi, F. Vitali (Eds.), *Knowledge Engineering and Knowledge Management, Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2016, pp. 239–254.