

Stimulating and Simulating Creativity with Dr Inventor

*Diarmuid P. O'Donoghue, *Yalemisew Abgaz, *Donny Hurley, §Francesco Ronzano, §Horacio Saggion

*Department of Computer Science, Maynooth University, Ireland.

§Universitat Pompeu Fabra, Barcelona, Spain
diarmuid.odonoghue@nuim.ie

Abstract

Dr Inventor is a system that is at once, a computational model of creative thinking and also a tool to ignite the creativity process among its users. Dr Inventor uncovers creative bisociations between semi-structured documents like academic papers, patent applications and psychology materials, by adopting a “big data” perspective to discover creative comparisons. The Dr Inventor system is described focusing on the transformation of this textual information into the graph-structure required by the creative cognitive model. Results are described using data from both psychological test materials and published research papers. The operation of Dr Inventor for both focused creativity and open ended creativity is also outlined.

Introduction

This paper describes the Dr Inventor project that is both a creativity support tool while its internal operation means that it also functions as a model of creative discovery. One of the core artifacts processed by Dr Inventor to boost scientific creativity is represented by Research Objects (RO) (Belhajjame et al., 2012), which are creative academic outputs including academic publications, patent applications and related data. Dr Inventor aims to actively explore creative bisociations (Koestler, 1964) between these Research Objects using a cognitively inspired model of creative thinking. This paper adopts a big data perspective on Research Objects attempting to uncover latent creative comparisons that might lie undiscovered within its dataset. Dr Inventor directly addresses two of Honavar's (2014) facets of computationally mediated scientific discovery: firstly the development of computational representations and secondly, computationally augmenting scientific discovery.

This paper is structured as follows. We first present a case for bisociative and analogy-based creativity, addressing some issues arising from Boden's attribution of bisociative reasoning to a category called “combinatorial creativity” (Boden, 1998). We then describe the Dr Inventor model, focusing on the processes that enable it to identify analogies between its text-based inputs. Next, we outline some results from text-based sources including human psychological tests and published research papers, illustrating

its operation as both a tool for focused creativity and also for open ended creativity. Finally a summary and some concluding remarks are made.

Analogical Reasoning and Creativity

The model of bisociative reasoning developed in this paper is built primarily on a computational model of analogical reasoning, which is extended to include additional background information. While computational treatments of analogy originally focused on the analogy per se, recent attention has focused more on situated models addressing topics like Ravens Progressive Matrices (Kunda, McGreggor and Goel, 2013). The analogy process provides a unique perspective from which to view computational creativity, lying at the crossroad of research in areas including cognitive science (Gick and Holyoak, 1980), developmental psychology (Rattermann and Gentner, 1998), computer science (Ramskar and Yarlett, 2003; O'Donoghue, Bohan and Keane, 2006; O'Donoghue and Keane, 2012) and neuroscience (Green et al., 2010). Research in these areas often constrain one another and offer the possibility of uncovering truly deep insights into the creative process. This may ultimately lead to formation of a cohesive multi-perspectival vision of one mode of creativity.

Analogy in Creative Reasoning

Psychological evidence has highlighted people's ability to reason using analogical comparisons in the laboratory setting (Gick and Holyoak, 1980). Subjects are typically presented with two analogous stories and are required to develop the latent analogy as a key to solving a problem in one of those stories. Later in this paper we shall demonstrate Dr Inventor's ability to take the texts used in these psychology tests and develop the same analogies as observed in (many) human participants in these trials.

The use of analogy has also been described in a “real world” scenario. Blanchette and Dunbar (2001) recorded and described the use of real-world analogies during laboratory meetings of molecular biologists and immunologists. They examined 16 different meetings in a number of different laboratories. They identified over 99 analogical comparisons and scientists typically used anything from 3

to 15 analogies in a one-hour meeting. The majority of the analogies discovered were between biological and immunological information – the so called “within-domains” analogies. However, the authors noted that scientists used more “between-domains” analogies (involving semantically distant source domains such as literature or engineering), when the goal involved a creative task such as formulating an hypothesis.

Goldschmidt *et al.*, (2011) and others have highlighted that “problem fixation” often frustrates peoples efforts to think creatively. That is, people experience difficulties in seeing new uses for existing information. The authors argue that to overcome this fixation and to promote creative thinking, that people be presented with semantically distant comparisons for a given problem. Research by Bowden *et al* (2005) and others has highlighted that insight occurs when problem solvers suddenly see a connection that previously eluded them. One possible mechanism of supporting insight is the discovery or a creative bisociations, like analogies and blends (Fauconnier and Turner, 1998).

Analogy and Transformational Creativity

Margaret Boden (Boden, 1990) offers three well-known levels of creativity, with increasingly impressive impact at the levels of *improbable*, *exploratory* or *transformational* creativity. Boden argues that analogy is effectively the lowest form of creativity (improbable); however we argue that when analogical reasoning is seen within the context of a cohesive system of human reasoning the picture is less clear. If the inferences mandated by an analogy contradict some fundamental axiomatic belief, especially beliefs with that numbers of associated deductions and inferences, then resolving this contradiction might well involve the “*shock and amazement*” associated with Boden's highest level of transformational creativity. It appears that analogies may in fact, drive creativity at any of Boden’s levels of creativity. Our creativity model is domain independent and does not include a pragmatic component or domain context. So, as our model does not use domain-specific knowledge, arguably it cannot be easily cast as one of improbable, exploratory or transformational creativity in Boden’s terms.

Creativity Producers and Consumers

Creativity is generally seen from the perspective of the creator. But, Dr Inventor needs to make a distinction between itself and its users who are consumers of its creative outputs. O’Donoghue and Keane (2012) made the point that a creative process may present a creative comparison so as to highlight the latent similarities, perhaps using terminology that highlights this commonality. However, discovering such creative comparisons will generally have to combat these differences in order to discover that commonality *ab initio*.

When they encounter a creative artefact, the interested consumer should also experience an episode of creativity, once they engage properly with the artefact. The process of engaging with a creative artefact should empower the consumer, ultimately leading them to a new conceptual space

akin to that of the creator. If the artefact doesn't cause this reaction, then its creative impact is greatly lessened and may be considered less creative. So, a truly creative output is not merely a recorded by-product of the creative experience of its creator, but it must also engender creativity within those consumers that engage properly with it. To achieve this, creative artefacts must have communicative potential and arguably, multiple creative artefacts may be necessary to clarify a new conceptual space - or to convince an unwilling consumer.

We call *secondary creativity* the act of engaging with a creative artefact so as to transform ones conceptual space, with *primary creativity* being the initial creative episode. We believe that secondary creativity is also essential for truly creative artefacts, helping wide adoption of this new perspective. Dr Inventor is concerned with both finding creative bisociations and with presenting these outputs to its users. It will use both ontology and visual analytics to support this secondary creativity.

Dr Inventor

Dr Inventor is a computationally creativity system that can both model scientific creativity and can also use its outputs to stimulate creative thinking within its users. It is as concerned with the process of creativity as it is with the products that arise from these processes (Stojanov and Indurkha, 2012). Dr Inventor is built on a cognitively inspired model of human bisociative reasoning, based on analogical comparisons and the counterpart projection of conceptual blends (Fauconnier and Turner, 1998; Veale, O’Donoghue and Keane, 2000). CrossBee (Jursic *et al.*, 2012) looked at exploring scientific papers, its focus lay in finding bridging terms between them. The focus of Dr Inventor is on finding and extending systemic similarities for creative purposes.

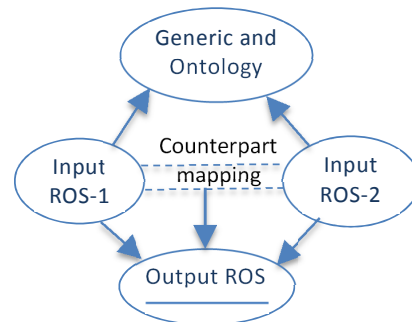


Figure 1: Conceptual spaces used by the bisociative model of Dr Inventor including the analogically founded mapping between the inputs

This paper focuses primarily on three of the four spaces of conceptual blending, namely the two input descriptions and the generic space. The dotted lines in Figure 1 indicate the correspondences between these inputs, derived with the help of Gentner’s structure mapping theory (1983). Dr Inventor’s 3-space model identifies a generic space contain-

ing the ontological similarity between paired relations from the Input 1 and Input 2. Dr Inventor thus identifies the generic space corresponding to the aligned items from the bisociation. This generic space also enables Dr Inventor to monitor the semantic congruity within a bisociation, to uncover comparisons more in fitting with the users' needs. Finally, the output space represents the new interpretation of one of those inputs. As each "target" maybe re-interpreted by multiple sources, and because that target may also act as a source for some other Research Object Skeleton (ROS), each newly created ROS is stored separately. For simplicity, this paper generally uses the terms *source* and *target*, unless specific point about the Blend is being made.

This means that a new ROS may act to later inspire subsequent creativity. Thus, Dr Inventor can potentially operate as a "Self-sustaining" creativity model as of described in O'Donoghue et al (2014). One of the chief obstacles hindering Dr Inventor in achieving this self-sustaining creativity lies in the quality of the new ROS and a sufficiently diverse knowledge base from which to progress.

The core data artefacts used by Dr Inventor are Research Objects (Belhajjame et al., 2012), which are research outputs including publications, patents, data, software (O'Donoghue et al., 2014b), social network information and other resources. Dealing with such heterogeneous data sources, characterized by consistent amounts of information to integrate and process, big data approaches and technologies are essential in order to enable the computational approaches to creativity in Dr Inventor. This paper focuses on the textual contents of RO, particularly of publications and patents. These documents are first subject to a number of processing activities to properly mine their contents in order to generate inputs that are useful to Dr Inventor's analogy-based model.

From each RO Dr Inventor generates a graph-based representation called the Research Object Skeleton (ROS) representing the key concepts and relationships extracted from that RO. Dr Inventor identifies similarities between these ROS with a view to extending these similarities and uncovering creative possibilities.

Dr Inventor Model

The overall Dr Inventor model contains components that deal with document summarization, information extraction; ontology learning, matching and personal recommendation; ROS generation, assessment, similarity and analogy/blending; validation, mapping, retrieval and finally visual analytics. The discussion in this paper will focus on the ROS generation, analogy/blending model and the creativity assessment components.

Mining Textual Contents to Populate ROS

In Dr Inventor, Research Object Skeletons (ROSs) are built by mining the contents aggregated by the corresponding Research Objects (ROs). To populate a ROS, Dr Inventor mainly relies on the extraction of information from the textual contents of a RO. To analyze these contents, Dr Inventor integrates a Natural Language Processing Pipeline

(DRI-NLP pipeline) that aggregates and customizes several Information Extraction (Piskorski and Yangarber, 2013) and Text Summarization (Saggion, 2014) approaches and tools.

Since scientific publications constitute one of the main kinds of textual documents included in a RO, DRI-NLP pipeline has been properly structured to support the analysis of research papers. The great majority of papers are currently available as PDF files. As a consequence, the conversion of PDF into plain text constitutes an essential prerequisite to properly perform any further text analysis. To this purpose, DRI-NLP pipeline relies on PDFX (Constantin, Pettifer and Voronkov, 2013) that converts a PDF document of a scientific publication to a semi-structured text (XML). The plain text output of PDFX is thus processed so as to identify sentences by means of a custom rule-based sentence splitter. Each sentence is processed by means of the MATE dependency parser (Bohnet, 2010) to extract dependency relations which are represented in a dependency tree. DRI-NLP pipeline dependency parser has been customized in order to properly deal with several peculiarities of scientific publications, including the presence of inline citations. In particular, inline citation markers like "(AuthorA *et al.*)" or "(1)" are excluded from the dependency tree if they have no syntactic functions in the sentence where they are present. Dr Inventor is focused on the discipline of computer graphics as its test-bed, thus a particular challenge has been dealing with the many mathematical expressions in these papers and allowing their treatment separately from the main body of the text. Besides dependency parsing, DRI-NLP pipeline enables the creation extractive summaries of papers by ranking their sentences by relevance (Saggion, 2014).




Figure 2: Processing PDF papers by Dr Inventor Natural Language Processing Pipeline

As result of dependency parsing, each word of a sentence is characterized by its Part-Of-Speech (POS) (noun, verb, adjective, etc.) and dependency relations (subject, object, verb chain, modifier of nominal, etc.). The linguistic information extracted from each publication can be condensed in the tables: the Syntactic dependency and the POS tag table. In particular, Figure 2 focuses on the analysis of a specific sentence taken from the abstract of a paper.

While Dr Inventor is focused on the test-bed of computer graphics publications, it remains a general model capable of dealing with arbitrary text inputs. This paper also uses data derived from psychology text materials and work is ongoing using the texts of patent applications.

ROS Generation

The next task for Dr Inventor is to generate a ROS from the results of the parsing process. The representation we chose for these graphs is sufficiently general to represent different types of RO. Since we want to structure objects and their inter-relationships this information is stored as a graph, aimed at supporting the later structure mapping process (Gentner, 1983).

<p style="text-align: center;">Karla the Hawk:</p> <p><i>Karla, an old hawk, lived at the top of a tall oak tree. One afternoon, she saw a hunter on the ground with a bow and some crude arrows that had no feathers. The hunter took aim and shot at the hawk but missed. Karla knew the hunter wanted her feathers so she glided down to the hunter and offered to give him a few. The hunter was so grateful that he pledged never to shoot at a hawk again. He went off and shot a deer instead.</i></p> <p style="text-align: center;">Zerdia True Analogy:</p> <p><i>Once there was a small country called Zerdia that learned to make the world's smartest computer. One day Zerdia was attacked by its warlike neighbor, Gagrach. But the missiles were badly aimed and the attack failed. The Zerdian government realized that Gagrach wanted Zerdian computers so it offered to sell some of its computers to the country. The government of Gagrach was very pleased. It promised never to attack Zerdia again.</i></p>

Table 1: Textual contents of “Karla” and “Zerdia”

Each ROS is constructed as an *attributed relational graph* (ARG), which is a directed graph where nodes and edges may contain additional properties like labels, categories and numeric values. If required, we can store additional identifying information (e.g. Author, Affiliation, etc.) within the graph, but this information is not required for the analogy process *per se*.

The primary information in a ROS is the concept nodes (nouns) and the relationships (verbs) between them. Concept nodes are not linked directly to one another but are connected with relation nodes. To generate the ROS we use the general structure “subject” – “verb” – object” - as required by SMT. These triples arise from the dependency and POS tables as the input to ROS generation. Early testing has shown that taking triples directly from the dependency table typically leaves many of them incomplete, leaving ROS without the necessary structure to support identification of creative inter-domain correspondences. Therefore, Dr Inventor performs a deeper exploration of the tables in order to generate more useable ROS structures.

By constructing a dependency graph from the tables and applying a set of heuristics to the graphs, a more complete set of triples is generated. The heuristics involve combining some of the nodes and tracing through the graph finding pairs for each verb.

Figure 3 depicts two ROSs generated for the “Zerdia” and “Karla” stories (Table 1) used in human psychological studies (Gentner and Landers, 1985). They were generated by the text mining and ROS generation techniques discussed earlier, but some manual post-editing was performed to identify co-referencing concepts nodes in the ROS. In the “Zerdia” story the word “it” is used twice, but the ROS were edited so one instance was replaced by the referent “Zerdia” and another by “Gagrach”. In the “Karla” story the word “she” refers to “Karla” and “he/him” refers to “hunter”. While these co-referents were resolved manually work is underway in the text pipeline to automatically resolve these referents.

Dr Inventor explicitly represents higher-order (causal relations connect first-order relations) relations within a ROS. A distinct set of nodes represents the higher-order relations, these connecting the first-order (and potentially other high-order) relations. However, ROS generated from within our Research Objects corpus show that high-order (causal) relations are rarely explicitly identified. As we shall see, this influenced our choice of mapping algorithm.

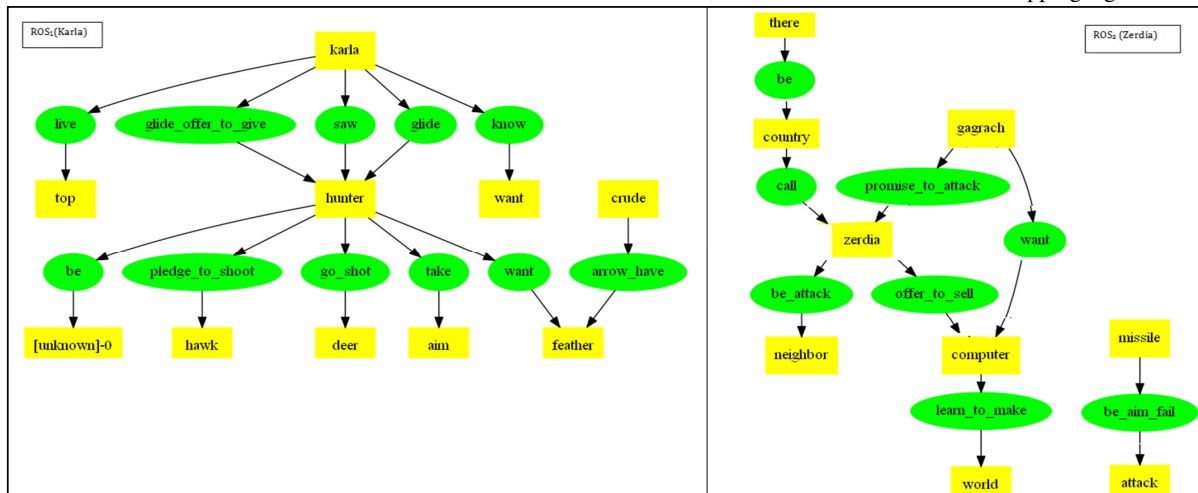


Figure 3: ROS for the “Karla” and “Zerdia” analogy used in human studies and Dr Inventor

Graph Storage

Dr Inventor uses the Neo4j graph database to store its ROSs. Neo4j has as its core structures; nodes, relationships between them and properties on both, this being the same structure as the ARG. Additional information such as the SentenceID or SectionTitle for each triple can also be stored in the Neo4j database. This can be useful when we want to map only between particular sections (e.g. Abstract or Conclusion) and also to reference back to the original sentences from which the triples were extracted.

Data Differences

Previous analogy models like SME (Forbus, Ferguson and Gentner, 1994), IAM (Keane and Bradshaw, 1988) and Kilaza (O'Donoghue and Keane, 2012) used hand coded data. The ROS generated above differs from the earlier hand-coded data in at least two significant respects. First ROS contain very few high-order relations, which are heavily used by mapping models mentioned above. Dr Inventor does not focus on the hierarchical structure of hand-coded data, using instead some lower level topological structure. Secondly (as mentioned in (O'Donoghue and Keane, 2012)) hand coded data often simplifies the comparison process by using relations that highlight the latent similarity. Dr Inventor must uncover and identify the hidden similarity even in the absence of such lexical cues.

Dr Inventors Creativity Engine

This paper focuses on the creativity engine that lies at the heart of Dr Inventor. Thus, we focus on creative analogy-based comparisons and show a number of features of Dr Inventor that specifically attempt to support the identification and generation of these creative analogies.

Creative Analogies

A number of properties appear to be shared amongst many creative analogical comparisons (O'Donoghue and Keane, 2012) and these facets are used to generate novel and potentially useful analogies and blends. Firstly the source (domain) of inspiration is typically semantically different from the given target problem. That is, creative sources tend to be sufficiently different and any similarity is non-obvious and has not been previously explored in detail. Secondly, the creative source contains the necessary structural similarity that is required to generations of viable analogy with the given problem.

To this end, Dr. Inventor specifically seeks out bisociations that involve two semantically distant domains, that form a rich inter-domain mapping and that yield inferences suggesting something new about one of those domains.

Graph Mapping

To support creative analogies Dr Inventor's retrieval and mapping activities makes frequent use of topological features derived from each ROS. For analogical mapping we exploit features such as type of the nodes (verb, noun), types of relationships (subject, object), degree (in-degree,

out-degree) and node rank values calculated by Node Rank algorithm (Bhattacharya et al., 2012).

We initiate the mapping process by calculating the Node Rank and by sorting the nodes in a descending order. The ranking allows us to start the graph matching process from the most centrally connected and useful node. This will further be used to serve as a threshold to screen useful nodes to improve the performance of the mapping process. The results presented in this paper have been generated using smaller RO (such as the abstracts of graphics paper RO), so performance has not been an issue. However this situation will change when mapping between ROS with large number of nodes is required.

The relation (verb) nodes in each ROS are represented distinctly, with one instance of a relation node for each verb contained in the RO. Verb nodes are central to the process of representing the content of the RO, however their connectedness is limited to a degree of 2 and thereby affects the resulting Node Rank values. However, multiple references to the same concept (noun) node will appear in the ROS as a single concept node – but referenced multiple times by each distinct relation node. Thus concept nodes have the greatest direct impact on the Node Rank values as a single concept node may be linked through many relations within a ROS. The mapping process avails of this referential structure when generating the largest graph-mapping between two ROS.

To identify a pair of mapping nodes from the source to the target, we used structural similarity score (using the connectedness of the nodes) and the literal similarity score. Using structural similarity, we consider two nodes as candidate mapping nodes if they have a higher similarity score. Whereas, literal similarity calculates the similarity coefficient between two words and yields a value between 0 and 1, where 0 indicates no similarity and 1 indicates complete similarity (synonym). This is achieved by using the Wu&Palmer (Wu and Palmer, 1994) WordNet-based similarity metric.

The mapping algorithm firstly selects a pair of nodes $P(sNode, tNode)$ from the source and the target respectively, with the highest node ranked nodes being selected first. In this way, the algorithm focuses on highly connected nodes within the graph because they contribute most to the mapping and analogical inference activities. Secondly, the mapping process checks if the selected pair $P(sNode, tNode)$ is structurally feasible for analogical mapping. A structurally feasible pair contains a source node which has degree (in degree and out degree) greater than or equal to degree (in degree and the out degree) of the target node respectively. The comparison ensures the identification of a sub-graph or an isomorphic graph of the target graph in the source graph. It further assesses the semantic similarity of the two nodes using Wu& Palmer. Next, mapping adds $P(sNode, tNode)$ to the inter-domain map to incrementally build a mapping sub-graph, if $P(sNode, tNode)$ is feasible. The mapping stores the pair of mapping nodes along with their similarity scores.

The mapping process then generates new candidate mappings by expanding *sNode* and *tNode* of $P(sNode, tNode)$ to their respective connected nodes that are not already expanded. By following the “subject” or “object” relationship path it reaches the connected nodes of the graph, incrementally adding these to the inter-domain mapping. After the candidate pairs are generated, they are ranked using their semantic similarity score. Ranking the candidate pairs will give us a chance to expand pairs with the highest semantic similarity first.

After including all mappings arising from the initial root mapping, the process then resumes with the next highest ranked and unmapped predicates. The algorithm employs depth first search to expand the nodes in the graph to identify new mapping pairs. Finally, it selects the mapping that contains the largest sub-graph and returns the mapping nodes together with their semantic similarity score.

We now look at the results produced by generating a mapping between the “Karla” and “Zerdia” psychology materials listed above, with the corresponding ROS being depicted in Figure 3. We note that this simulation of human analogy process began with the same text materials that were presented to human subjects. This comparison is an example to focused creativity, where both the source and target have been pre-identified.

The mapping between “Karla” and “Zerdia” gives us 11 mapped nodes between the source and target (Table 2). For example the noun node “Karla” maps to “Zerdia”, “Feather” maps to “Computer” and “Hunter” maps to “Gagrach”. Such a mapping identifies analogous items between the source and the target and is crucial for transferring new knowledge from the source to the target. In this specific example 50% of the nodes in the target ROS are mapped to the source ROS with an average Wu&Palmer similarity score of 0.56. The original domains can often include information that does not participate in the mapping, such as the (*missile be-take-aim attack*) in the Zerdia story. However the absence of this relation from the mapping is not terribly significant as it is an isolated fragment of information and does not contribute largely to the main story – that contributes to the largest connected component of that ROS.

Mapping Nouns		Mapping Verbs	
Source	Target	Source	Target
Hunter	Gagrach	Want	Want
Crude	World	Live	Offer_to_sell
Feather	Computer	Arrow_have	Learn_to_make
Want	Country	Glide_offer_to_give	Be_attack
Karla	Zerdia	Glide	Promise_to_attack
		Know	Call

Table 2: Mapping between “Karla” and “Zerdia”.

Analogies within Graphics Collections

To examine the mapping process, we used 10 papers from computer graphics domain. The abstracts of the papers were extracted and were processed using the previously described steps. Each ROS were mapped to the other 10

ROS including itself. The most basic step is to compare a ROS against itself. For all the 10 papers Dr Inventor yields the highest number of mapped nodes when a ROS is compared with itself – with all or almost all nodes being successfully mapped. This could be considered as a very basic step toward the evaluation of the mapping component of Dr Inventor.

The mapping of a ROS against the remaining 9 ROS identifies pairs that have the highest mapping nodes and pairs that have the lowest mapping nodes. The most analogous papers are those with large number of mapped nodes and highest similarity score. For example, the most similar non-identical mapping among the 10 papers is between “*Bar-Net_Driven Skinning for Character Animation*” and “*Real Time Large Deformation Character Skinning in Hardware*” with 14 mapping nodes and an average Wu&Palmer similarity score of 0.36. While the semantic similarity score may appear quite low, this was achieved from within a small collection of papers. We conducted a quick manual comparison between the abstracts of these papers and initial results indicate that these papers can be considered somewhat analogous to one another as for example, both papers present different approaches to the computer graphic topics of “skinning”. This analogy arose from the desire to identifying the largest mapping with the strongest semantic similarity from within the 10 papers however, the next section will discuss a more creative Use Case scenario.

The lowest mapping occurs between the papers “*Curve Skeleton Skinning for Human and Creature Character*” and “*Pose Space Deformation: A Unified Approach to Shape Interpolation and Skeleton-Driven Deformation*” with 5 mapping nodes and average similarity of 0.35. The mapping process, as it is expected, is not symmetrical, i.e. mapping between (*S, T*) may not be the same as mapping between (*T, S*). For example, saying “*a man is like a pig*” is not the same as saying that “*a pig is like a man*”! However, in this specific data set, it does not significantly affect the highest mapping node pairs.

Analogical Inference and Pattern Completion

Once we find a mapping between the two ROS, the next phase is to generate the resulting inferences by applying a “pattern completion” process to that mapping. This adds the newly inferred information to the target ROS to produce the new interpretation of that concept. In the exploratory analogy mapping process, the user may be interested to explore all the candidate nodes once he/she knows the existence of analogy between the source and the target.

Creativity Support and Evaluation in Dr Inventor

Dr Inventor is focused firstly on operating as a Creativity Support Tool (CST) and secondly, as a simulation of the analogy process. Shneiderman (2007) noted that there are no obvious metrics to quantify for CST’s and this problem lies at the core of creativity assessment and evaluation. The following two approaches are useful to evaluate the level of creativity support provided by Dr Inventor.

Among the functionality being developed for Dr Inventor is an “Inspire Me” Use Case, enabling users to creatively re-interpret one of their own papers. This will be achieved by using the paper as a target and searching the archive for papers that can act as a creative source domain, forming a large and semantically well-balanced mapping by making use of the topological structure of each ROS. Dr Inventor will identify and present to the user those analogies offering a collection of novel inferences that highlight the potential benefits of adopting this creative analogical comparison. Internal metrics will serve to select the most promising analogies to present to users, assessing the structural and semantic foundations of the comparison.

Implicit feedback on the presented analogies will be gathered by the user interface, enabling comparative evaluation of different comparisons by monitoring user engagement. Explicit user feedback also plays a very important role in evaluating Dr Inventor, using experts in computer graphics for evaluation. The Creative Support Index (CSI) (Cherry and Latulipe, 2014) is a psychometric survey that will serve to assess the creativity support provided by Dr Inventor. It is quick and easy to administer and is composed of two sections; a rating scale section and a paired-factor comparison section. It identifies 6 major factors of creativity, namely: enjoyment, exploration, expressiveness, immersion, results worth effort and collaboration. Under each of the factors, CSI asks two questions that are rated between 0 and 10, where 0 indicates the lowest value and 10 indicate the highest achievement. The paired-factor comparison section consists of each factor paired against every other factor for a total of 15 comparisons. As Dr Inventor will support users with different levels of expertise (first year PhD students to experienced Professors), this factor in particular will have to be controlled monitored during the evaluation process.

The Creative Achievement Questionnaire (CAQ) (Carson, Peterson and Higgins, 2005) is a very broad and general creativity assessment technique. Within the context of Dr Inventor, achievements appear to be primarily assessed by qualifying the number of published scientific papers. But the CAQ provides poor coverage of lower levels of creative achievement (before publication) that could guide development of the Dr Inventor project. However, the CAQ might be useful for the final evaluation of Dr Inventor.

We also note that (Jordanous, 2014) identifies five criteria to support meta-evaluation of computational creativity *per se*, as opposed to our current focus on Dr Inventor as a creativity support tool.

Summary and Conclusions

Dr Inventor is a computationally creative model that acts as both a simulation of human creative reasoning and also as a creativity support tool. We described how Dr Inventor performs text extraction from research publications presented in pdf format, describing how it addresses many complications that result from use of the pdf format. The dependency parser was described, as was the process of

constructing the graph representation used by the core model. Some peculiarities of the resulting graphs were noted, particularly the extreme rarity of identifiable higher-order causal relations. Some implications were noted the process of identifying the inter-domain correspondence. The text used from human psychological trials showed the ability of Dr Inventor to generate comparison using these same textual materials. Results for other Research Objects were outlined.

The mapping and evaluation process uses ontological information as a preference criterion to choose comparisons with the greatest potential for creativity on Dr Inventor users. Ontology also opens the way to re-describe the original documents, highlight the identified similarities. While this work is ongoing, it opens the way for early evaluation of Dr Inventor by comparing the impact upon users of creatively re-interpreted documents.

Acknowledgements

The research leading to these results has received funding from the European Union Seventh Framework Programme ((FP7/2007-2013)) under grant agreement no 611383.

References

- Belhajjame, K., Corcho, O., Garijo, D., Zhao, J., Missier, P., Newman, D., Bechhofer, S., García, E., Gómez-Pérez, J.M., Palma, R., Soiland-Reyes, S., Verdes-Montenegro, L., Roure, D.D. and Goble, C. (2012) 'Workflow-centric research objects: First class citizens in scholarly discourse', 9th Extended Semantic Web Conference, Hersonissos, Greece.
- Bhattacharya, P., Iliofotou, M., Neamtii, I. and Faloutsos, M. (2012) 'Graph-based Analysis and Prediction for Software Evolution', in *Proceedings of the 34th International Conference on Software Engineering*, Piscataway, NJ, USA: IEEE Press.
- Blanchette, I. and Dunbar, K. (2001) 'Analogy use in naturalistic settings: The influence of audience, emotion, and goals', *Memory & Cognition*, vol. 29, no. 5, pp. 730-735.
- Boden, M. (1990) *The Creative Mind*, Weidenfeld and Nicolson.
- Boden, M. (1998) 'Computer Models of Creativity', in Sternberg, R.J. *Handbook of Creativity*, Cambridge University Press.
- Bohnet, B. (2010) 'Very high accuracy and fast dependency parsing is not a contradiction.', Proc. 23rd International Conference on Computational Linguistics, 89–97.
- Bowden, E.M., Jung-Beeman, M., Fleck, J. and Kounios, J. (2005) 'New approaches to demystifying insight', *Trends in cognitive sciences*, vol. 9, no. 7, pp. 322-328.
- Brown, T.L. (2003) *Making Truth: Metaphor in Science*, University of Illinois Press.
- Carson, S., Peterson, J.B. and Higgins, D.M. (2005) 'Reliability, Validity, and Factor Structure of the Creative

- Achievement Questionnaire', *Creativity Research Journal*, vol. 17, no. 1, pp. 37–50.
- Cherry, E. and Latulipe, C. (2014) 'Quantifying the Creativity Support of Digital Tools through the Creativity Support Index', *ACM Transactions on Computer-Human Interaction*, vol. 21, no. 4.
- Constantin, A., Pettifer, S. and Voronkov, A. (2013) 'PDFX: fully-automated PDF-to-XML conversion of scientific literature', Proceedings of the 2013 ACM symposium on Document engineering, 177–180.
- Fauconnier, G. and Turner, M. (1998) 'Conceptual integration networks', *Cognitive science*, vol. 22, no. 2, pp. 133–187.
- Forbus, K.D., Ferguson, R.W. and Gentner, D. (1994) 'Incremental structure-mapping', Proc 16th Conference of the Cognitive Science Society, 313–318.
- Gentner, D. (1983) 'Structure-mapping: A theoretical framework for analogy', *Cognitive Science*, vol. 7, pp. 155–170.
- Gentner, D. and Landers, R. (1985) 'Analogical reminding: A good match is hard to find', Proceedings of the International Conference on Systems, Man, and Cybernetics., Tucson, AZ.
- Gick, M.L. and Holyoak, K.J. (1980) 'Analogical problem solving', *Cognitive psychology*, vol. 12, no. 3, pp. 306–355.
- Goldschmidt, G. (2011) 'Avoiding Design Fixation: Transformation and Abstraction in Mapping from Source to Target', *The Journal of Creative Behavior*, vol. 45, no. 2, pp. 92–100.
- Green, A.E., Kraemer, D., Fugelsang, J.A., Gray, J.R. and Dunbar, K.N. (2010) 'Connecting long distance: semantic distance in analogical reasoning modulates frontopolar cortex activity', *Cerebral Cortex*, vol. 20, no. 1, pp. 70–76.
- Honavar, V.G. (2014) 'The Promise and Potential of Big Data: A Case for Discovery Informatics', *Review of Policy Research*, vol. 31, no. 4, pp. 326–330.
- Jordanous, A. (2014) 'Stepping Back to Progress Forwards: Setting Standards for Meta-Evaluation of Computational Creativity', Proc ICCS, 129–136.
- Jursic, M., Cestnik, B., Urbancic, T. and Larvac, N. (2012) 'Cross-domain Literature Mining', Proc ICCS, pp 33–40.
- Keane, M.T. and Bradshaw, M. (1988) 'The Incremental Analogical machine', in Sleeman, D. (ed.) *3rd European Working Session on Machine Learning*, Kaufmann, CA, USA.
- Keane, M.T., Ledgeway, T. and Duff, S. (1994) 'Constraints on Analogical Mapping: A Comparison of Three Models', *Cognitive Science*, vol. 18, no. 3, pp. 387–438.
- Koestler, A. (1964) *The Act of Creation*, Penguin Books.
- Kuhn, T.S. (1962) *The structure of scientific revolutions*, University of Chicago press.
- Kunda, M., McGreggor, k. and Goel, A.k. (2013) 'A computational model for solving problems from the Raven's Progressive Matrices intelligence test using iconic visual representations', *Cognitive Systems Research*, vol. 22–23, pp. 47–66.
- O'Donoghue, D.P., Bohan, A. and Keane, M.T. (2006) 'Seeing Things: Inventive Reasoning with Geometric Analogies and Topographic Maps', *New Generation Computing*, 24, (3) (special issue on Computational Creativity), pp. 267–288.
- O'Donoghue, D.P. and Keane, M.T. (2012) 'A Creative Analogy Machine: Results and Challenges', 4th International Conference on Computational Creativity (ICCC), UCD, Dublin, Ireland, 17–24.
- O'Donoghue, D.P., Monahan, R., Grijincu, D., Pitu, M., Halim, F., Rahman, F., Abgaz, Y. and Hurley, D. (2014b) 'Creating Formal Specifications with Analogical Reasoning', PICS-Publication Series of the Institute of Cognitive Science.
- O'Donoghue, D.P., Power, J., O'Briain, S., Dong, F., Mooney, A., Hurley, D., Abgaz, Y. and Markham, C. (2014) 'Can a Computationally Creative System Create Itself? Creative Artefacts and Creative Processes', International Conference on Computational Creativity (ICCC), Ljubljana, Slovenia.
- Piskorski, J. and Yangarber, R. (2013) 'Information extraction: Past, present and future', *Multi-source, multilingual information extraction and summarization*, pp. 23–49.
- Ramsar, M. and Yarlett, D. (2003) 'Semantic grounding in models of analogy: an environmental approach', *Cognitive Science* 27, pp. 41–71.
- Rattermann, M.J. and Gentner, D. (1998) 'More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task', *Cognitive Development*, vol. 13, no. 4, pp. 453–478.
- Saggion, H. (2014) 'Creating Summarization Systems with SUMMA.', Proceedings of Language Resources and Evaluation Conference.
- Shneiderman, B. (2007) 'Creativity support tools: Accelerating discovery and innovation', *Communications of the ACM*, vol. 50, no. 12, pp. 20–32.
- Silvia, P.J., Wigert, B., Reiter-Palmon, R. and Kaufman, J.C. (2012) 'Assessing Creativity with Self-Report Scales: A Review and Empirical Evaluation.', *Psychology of Aesthetics, Creativity, and the Arts*, vol. 6, no. 1, pp. 19–34.
- Stojanov, S. and Indurkha, B. (2012) 'Perceptual Similarity and Analogy in Creativity and Cognitive Development.', SAMAI Workshop at ECAI, Montpellier France, 19–24.
- Veale, T., O'Donoghue, D.P. and Keane, M.T. (2000) 'Computation and Blending', *Cognitive Linguistics* 11 (3/4), pp. 253–281.
- Wu, Z. and Palmer, M. (1994) 'Verbs Semantics and Lexical Selection', in *Proceedings of the 32nd Annual Meeting on Association for Computational Linguistics*, Las Cruces, New Mexico: Association for Computational Linguistics.