The MindSpaces Knowledge Graph: Applied Logic and Semantics on Indoor and Urban Adaptive Design

Evangelos A. Stathopoulos[1] [a], Alexandros Vassiliades[1,2] [b], Sotiris Diplaris[1] [c], Stefanos Vrochidis[1] [d], Nick Bassiliades[2] [e] and Ioannis Kompatsiaris[1] [f]

[1] *Information Technologies Institute, Center for Research and Technology Hellas, Thessaloniki, Greece*
[2] *School of Informatics, Aristotle University of Thessaloniki, Greece*

Keywords: Knowledge Graph, Ontology, Reasoning Mechanism, Indoor/Urban Adaptive Design, Semantic Web.

Abstract: The evolution of Knowledge Graphs (KGs), during the last two decades, has encouraged developers to create more and more context related KGs. This advance is extremely important because Artificial Intelligence (AI) applications can access open domain specific information in a semantically rich, machine understandable format. In this paper, we present the MindSpaces KG, a KG that can represent emotions-relevant and functional design for the indoor and urban adaptive design. The MindSpaces KG can integrate emotional, physiological, visual, and textual measurements, for the development of online adapting environments. Moreover, we present a reasoning mechanism that extracts crucial knowledge from the MindSpaces KG, which can help users in real-life scenarios. The scenarios were provided by experts.

1 INTRODUCTION

The evolution of Knowledge Graphs (KGs) in the last 20 years allowed developers to construct context related KGs (i.e., KGs that can be used only in specific environments). The creation of context related KGs seems to be the next step for allowing KGs to become the main knowledge representation format for the Web (Berners-Lee et al., 2001). Our focus is on representing emotions-relevant and functional design for the indoor and urban adaptive design. Emotions-relevant refer to emotions created to individuals when experiencing an indoor or urban area, such as stress, calmness, happiness, among others. The functional design of an indoor or urban location imply the practicality of the location, for example movability of workers in a workspace in the case of indoor environments, or redesign-relocate of a non helpful bus station in the case of urban environments.

The MindSpaces KG was developed in order to work as the knowledge representation of the MindSpaces project[1]. The motivation for the MindSpaces KG stands in 2 different key points. The first is to improve urban design in a rapidly expanding city by addressing new challenges that may arise related to its functionality, mobility attractiveness, protection of culture and environment. MindSpaces KG will increase sensitivity and awareness towards the cultural significance and current issues of a city, related to the environment and mobility. While the second, is to increase opportunities for positive social interaction in work environments which leads to improved productivity and creativity across departments and teams, by helping to readjust workspaces to achieve better aesthetics and functionality.

The MindSpaces KG is mostly oriented for artists, designers and architects for redesigning indoor workspaces in order for the workers to feel more comfortable, and to improve functionality, mobility and overall attractiveness of areas. But other users can also utilize the MindSpaces KG, as for instance the ideas of citizens of a city can be very useful in redesigning urban areas. Moreover, we provide a reasoning mechanism that aids users in real-life scenarios which were provided by experts.

The problem we are addressing is the construction of a general KG, which can represent emotions-

[a] https://orcid.org/0000-0003-3713-5833
[b] https://orcid.org/0000-0003-4569-503X
[c] https://orcid.org/0000-0002-9969-6436
[d] https://orcid.org/0000-0002-2505-9178
[e] https://orcid.org/0000-0001-6035-1038
[f] https://orcid.org/0000-0001-6447-9020

[1] https://mindspaces.eu/

Stathopoulos, E., Vassiliades, A., Diplaris, S., Vrochidis, S., Bassiliades, N. and Kompatsiaris, I.
The MindSpaces Knowledge Graph: Applied Logic and Semantics on Indoor and Urban Adaptive Design.
DOI: 10.5220/0011666600003393
In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART 2023) - Volume 3*, pages 334-341
ISBN: 978-989-758-623-1; ISSN: 2184-433X

relevant and functional design for indoor and urban adaptive design, by mapping emotional, physiological, visual, and textual measurements. Moreover, we address the fact that a user in order to access the information inside the KG in a convenient way, must be provided with a reasoning mechanism based on real-life scenarios. The real-life scenarios ideally should be formulated by domain experts.

The scalability of the MindSpaces KG is wide, as it can represent information for any urban area and indoor workspace. More specifically, for the indoor workspaces the MindSpaces KG is mostly oriented for offices, but also other indoor workspaces, such as a cottage industry, and can also exploit the information inside the MindSpaces KG. A natural extension of it could be to extend representing emotions relevant to outdoor workspaces, and recommend tasks to citizens on how they can increase the functionality, mobility, and attractiveness of an urban area with respect to the cultural significance and environment.

The problem tackled with this semantic framework has to do with the different nature of data present in the system. Several approaches can be found that address the knowledge representation with ontologies in each domain distinctly. The motivation behind this work is that unique semantic requirements needed to be addressed altogether and combine interdisciplinary domains in one unique ontology, tied together carefully with precise and meaningful custom interconnections to serve multimodal knowledge representation and enable smart reasoning mechanisms.

Our contribution with this paper, is on one hand the MindSpaces KG which can represent multi-modal measurements, which in turn help artists with the empirical and pragmatic perception of actual occupants, so as to drive the development of unconventional solutions in the design of spaces. On the other hand, is the reasoning mechanism for the MindSpaces KG, which proves to be helpful in real-life scenarios.

The rest of this paper is organized as follows. Section 2, contains the related work. Next, in Section 3 we present the MindSpaces KG, the reasoning mechanism, and the data upon we constructed the KG. Section 4, contains the evaluation of the KG and the reasoning mechanism. We conclude our paper with Section 5.

2 RELATED WORK

The study has two main directions, the MindSpaces KG, and the reasoning mechanism which extracts knowledge from the MindSpaces KG. For this reason, we will separate our related work into two main subsections, one for similar knowledge graphs, and one for the domain specific reasoning mechanisms from knowledge graphs.

Knowledge Graphs: The first category of knowledge graphs that can be considered close to the MindSpaces KG, are KGs for arts and artists. These KGs fall into the category of KGs for cultural heritage (Hyvönen, 2012; Schneider, 2020). Moreover, the area of KGs about arts and artists is not so rich, as there are not many studies that could be classified clearly in this domain. One exception, is the study of Raven et al. (Raven et al., 2020), where the authors present a KG that can represent the steps of specification, conceptualization, integration, implementation and evaluation in a case study concerning ceramic-glass. The difference with MindSpaces KG, can be noticed, as we offer a KG focused on the knowledge representation of emotions-relevant and functional design for the indoor and urban adaptive design.

Thereinafter, the area of KGs for architecture and architects seems to be quite richer. (Lopes, 2007), showcases the main notions that must be represented in a KG about architectural concepts. On the other hand, KGs like (Wagner and Rüppel, 2019; Kumar et al., 2019), which are considered in the area of architecture mostly contain information about materials of objects and their uses. For the former, we noticed it is a theoretical study, while we offer a constructed KG. For the latter three KGs, our KG does not contain similar type of information, as we offer knowledge about different interdisciplinary distinct domains.

Reasoning: The area of reasoning over KGs, is quite rich and it was enhanced in the last decade with the constant evolution of KGs. The reason why KGs are very helpful in retrieving information, is because they have a predefined format, which is easily understandable by the machine, and their terminology and assertion components can be defined based on a set of rules which can represent commonsense knowledge. Even though many general techniques have been presented (Munir and Anjum, 2018; Asim et al., 2019; Yu, 2019), there is lack of studies for reasoning over KGs restricted to a specific domain. The reason why is hard to create a reasoning mechanism over KGs which are restricted to a specific domain, is because domain experts are needed in order to create real-life scenarios, based on which the reasoning mechanism should be constructed (as mentioned in (Chi et al., 2019)). (Vassiliades et al., 2020) is an exception as a reasoning mechanism for the household environment is presented, but in contrast to our study the scenarios presented are not a result of knowledge provided by domain experts.

3 THE MindSpaces KNOWLEDGE GRAPH

The MindSpaces KG is part of the MindSpaces project. Therefore, the MindSpaces KG communicates with other components of the project. We mention this because some parts of the mapping mechanism, which receives messages, in JSON format, from the visual, textual, and stress analysis component will not be analyzed in detail in this paper. But the source code of the mapping mechanism can be found here[2]. The idea of the pipeline is that after the mapping mechanism has received the message from a component, it will map the information into the KG. Then, any user can hand-pick from a set of predefined SPARQL queries, in order to access the information in the KG. Notice that all SPARQL queries were defined with the help of domain experts, and some of them also require an image as an input in order to cast the SPARQL query. Figure 1, shows an outline of the pipeline, each number in the circles shows the order of steps.

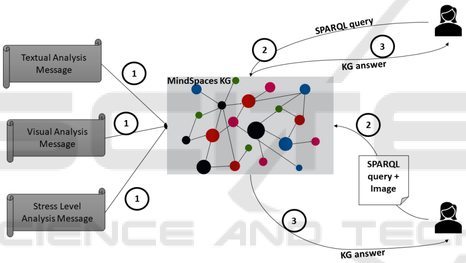

Figure 1: Pipeline of the MindSpaces KG.

3.1 Nature of Data

The multi-modality and variety of data flowing in the system and the necessity of homogenization and fusion mandated the adoption of a semantic knowledge graph to address the requirements of the project. The knowledge graph is not responsible for archiving and storing raw data files, since there is an underlying data storage facility for that purpose. Instead, the KG hosts metadata of raw data, analyses results and miscellaneous information with semantic value among other candidates for being mapped and fused into the knowledge base. One can find a blueprint of the messages fused in the KG here[2]. The KG is accompanied with novel ontological models to achieve proper semantic annotation of the raw data.

The main categories of data which needed to be captured in the knowledge graph were: physiological analysis results from galvanic skin responses (GSR), visual analysis results from images, textual analysis results derived from online retrieved content, and general information about VR experiments. Some of the aforementioned analyses results share some common relations such as in the case of the imageability metric.

The physiological signals (GSR) were captured during hot-spot or navigation VR experiments where multiple users conducted stress-induced tasks inside virtual environments containing multiple configurations of work offices. The ultimate goal was to obtain and assess stress indicators directly from the skin of the subjects while experiencing different setups of the space.

The visual analysis component consists of several machine learning models either trained from scratch or fine-tuned from other efforts. The model performing semantic segmentation (Qiu et al., 2021) on images was trained to extract semantic labels and percentages per pixel on images while the Verge classifier (Andreadis et al., 2020) was deployed to classify the images to one or more classes based on context. The valence-arousal model for exterior design was trained on a newly collected dataset annotated by experts to deliver confidence and values on a happy-unhappy and calm-excited scale, while a visual imageability score is generated (Pistola et al., 2022).

3.2 The MindSpaces Knowledge Graph

The MindSpaces KG is separated into two big sub graphs. One for the purpose of representing knowledge that aims at improving urban design in a rapidly expanding city by addressing new challenges that may arise related to its functionality, mobility, attractiveness, protection of culture and environment. The other aims at representing knowledge that involves increasing opportunities for positive social interaction in work environments which leads to improved productivity and creativity across departments and teams, by helping to readjust workspaces to achieve better aesthetics and functionality.

The subgraph for the first case was constructed based on the information we received from the various components. In more detail, we analysed the messages that we received from the textual and visual analysis components, and we developed a KG based on those. In the corresponding part of the ontology we represent the important concepts of these messages through classes and relations among them. Moreover, we analysed the requirements of the reasoning mechanism, meaning that we took into consideration the competency questions given by users and experts, and defined classes and relations in such a

[2]https://github.com/valexande/MindSpacesPUC1-2

way that would help the reasoning mechanism returning the crucial information. Notice that a competency question, is a question which as a user we would like to be answered by the KG, with the information that it contains. In Figure 2 one can see the classes and the object type properties of the first subgraph.

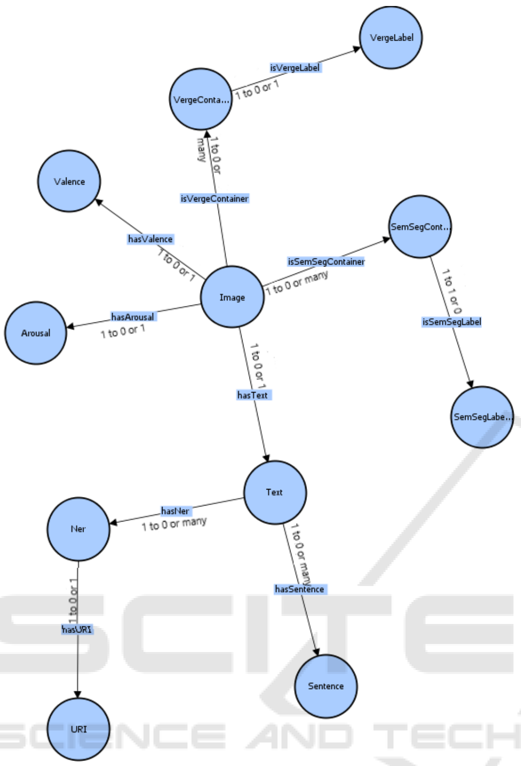

Figure 2: The schema of the first subgraph of the MindSpaces KG.

The namespace *mind1* is used to indicate the classes and relations for the first subgraph of the MindSpaces KG. Next, we will give a detailed analysis of the classes and the relations between them.

- The **Sentence** class contains information about the sentences that compose a textual description of an image. For each sentence it has information about the emotional tag that was given by the user, and the confidence of each emotional tag. The emotional tag is the sentiment label that was given by the user to a sentence, such as positive, negative, and neutral.
- The **Text** class contains information about the textual description of an image. The Text class is connected through the property hasSentence with the class Sentence, in order to give further information about the sentences that compose the textual description. Moreover, it has information for the language of the textual description, and the textual description itself. Finally, the Text class is connected through the property hasNer with the Ner class, which has information about named entities found in the textual description. Named entities can be words that refer to real life objects, actions, or activities.
- The **Ner** class contains information about the named entities found in the textual description of an image. The Ner class gives information for the category of the named entity relation (the category of a named entity is a classification that was given by the domain experts), the imageability score of the named entity, and is connected through the property hasURI with the class URI. The imageability score is a confidence score that is composed by the visual analysis component. Imageability according to the urban planner Kevin Lynch is the quality of a physical object to evoke a strong image in any observer, thus being memorable (Lynch et al., 1960).
- The **URI** class contains information about the URIs of the named entities. Currently, the URIs point only to DBpedia entities. The URI class has information about the URI link and the confidence that a named entity should be related with a specific URI.
- The **VergeLabel** class contains information about the labels of the verge classifications found in an image.
- The **VergeContainer** class contains information about the imageability scores of the verge classifications found in an image. Moreover, it is connected though the property **isVergeLabel** with the class **VergeLabel**, in order to indicate the label of a verge classification.
- The **SemSegLabel** class contains information about the labels of the segmented objects found in an image.
- The **SemSegContainer** class contains information about the imageability scores of the segmented objects found in an image, the percentage of space they capture in the image, and the confidence that they are part of the image. Moreover, it is connected though the property **hasSemSegLabel** with the class **SemSegLabel**, in order to indicate the label of a segmented object.
- The **Arousal** class contains information about the arousal score that was given by a user for an image. Arousal is a confidence score given by the users.
- The**Valence** class contains information about the valence score that was given by the user for an

image. Valence is a confidence score given by the users.

- The **Image** class is the most important class, as it contains a lot of metadata information about the characteristics of the image, such as the latitude, longitude, the pitch, the zoom, and others. But apart from these it is connected: (i) with the Arousal class through the property **hasArousal** to indicate its arousal, (ii) with the **Valence** class through the property **hasValence** to indicate its valence, (iii) with the **VergeContainer** class through the property **isVergeContainer** to give information about the verge classifications that it has, (iv) with the **SemSegContainer** class through the property **isSemSegContainer** to give information about the segmented objects it contains, and (v) with the **Text** class through the property **hasText** to give information about the textual description that it has.

In Figure 3 one can see the classes and the object type properties of the second subgraph.

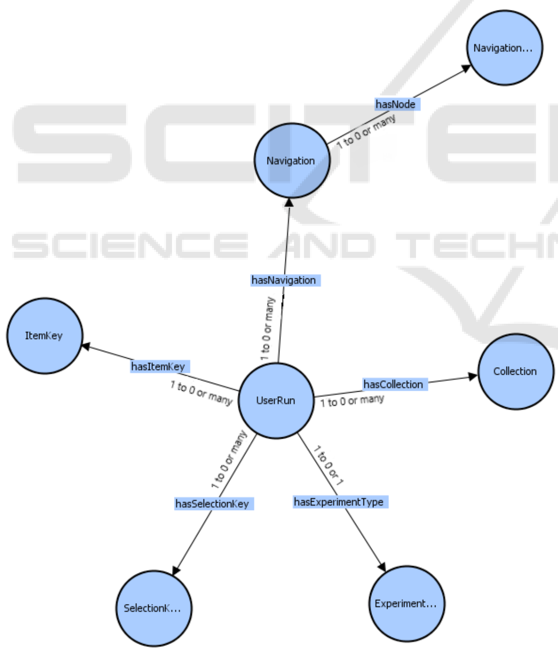

Figure 3: The schema of the second subgraph of the MindSpaces KG.

The namespace *mind2* is used for the second subgraph. Next, follows a detailed analysis of the classes and the relations of the second subgraph.

- The **Collection** class contains information about the experiment status data change, meaning that it indicates when the user run has started, stopped, and if it goes from state ON to state OFF.

- The **ExperimentType** class contains information about the experiment type data changes, meaning that it indicates if the experiment type is a navigation task, a navigation selection, hot spot experiment, or it does not have a type.

- The **Navigation** class contains information about the navigation configuration, meaning that it contains information about the type of the configuration, and the timestamp that the node was captured. Moreover, it is related through the property **hasNode** with the class **NavigationNode** that contains information about the navigation instances (i.e., navigation nodes).

- The **NavigationNode** class contains information about the navigation nodes, such as the x, y, z coordinates of the node, the fusion score, and the GSR score. The fusion and the GSR score, are some confidence score given by the visual and stress level analysis components.

- The **ItemKey** class contains information about the item key information, such as item key label, the collaboration, focus work, overall design, privacy, and stress scores for each item key. All of the aforementioned scores, are some confidence scores given by the users.

- Finally, the **UserRun** class is the most important class, as it connects the information from the aforementioned classes with a user run. More specifically the **UserRun** class is connected: (i) with the **ExperimentType** class through the property **hasExperimentType** to indicate the experiment types that it contains, (ii) with the **ItemKey** class through the property **hasItemKey** to indicate the item keys that it contains, (iii) with the **Collection** class through the property **hasCollection** to give information about the collections that it contains, and (iv) with the **Navigation** class through the property **hasNavigation** to give information for navigation nodes that it contains.

Notice that the two subgraphs eventhough they refer to different notions, i.e., the first subgraph is for the outdoor adaptive design and the second for the redesign of internal workspaces, they are related between them with the property **hasImage**. The property hasImage has domain the class **Collection** and range the class **Image**. The reason for that is because a collection is a set of images.

3.3 Reasoning and Logic

The reasoning mechanisms take advantage of the MindSpaces KG which was created to support the use cases of urban environments and indoor workspaces, as well as of the population of the knowledge base with content and metadata deriving from both artists and users.

The main idea, for the reasoning over the first subgraph (see subsection 3.2), was to feed on demand an interactive google map with geolocated 2D points in the form of CSV files corresponding to image entries inside the knowledge base through application programming interfaces (APIs) developed in node.js and Javascript. The graphdh.js library was used to establish the connection and transactions towards and from the GraphDB repository[3] with authorization and authentication ensured. The file delivered in response follows a scalable and dynamic approach, meaning it is being generated on demand based on live requests, thus ensuring always up-to-date data delivery as the knowledge base supports a continuous online population and new entries may arrive anytime. The format of the file consists of as many lines as the images fulfil the SPARQL queries and 5 columns:

- latitude (the latitude when the corresponding image was captured)
- longitude (the longitude when the corresponding image was captured)
- point size (the size of the circle to be depicted on the map)
- point opacity (the opacity of the circle to be depicted on the map)
- point colour (the colour of the circle to be depicted on the map)

In total, 16 SPARQL queries were formulated, each satisfying a different user requirement, followed by an additional multipurpose sparse function. But due to space restrictions only two Scenario A and Scenario B, which were defined by domain experts as real-life scenarios, will be analyzed.

Scenario A: Given a snap image, return a list of images, where: (i) the Top 3 segmentation labels of the snap (based on the coverage percentage) exist in the images, (ii) the images must have imageability $\geq$ imageability of snap + 0.05, (iii) the Top 3 segmentation labels must exist in the images with the same coverage percentage, or a 20% difference, and (iv) the results must be limited to 8 images, if there exists as many, sorted based on their imageability.

Scenario B: Given a snap image, return a list of images, where: (i) the Top 3 segmentation labels of the snap (based on the colorfulness percentage) exist in the images, and their colorfulness is above 1, and (ii) for each one of the 3 segmentation classes bring the Top 5 images with the highest colorfulness for each segmentation classes.

Notice that the queries which are refered in this subsection are also Competency Questions (CQs), except from **Scenario A** and **Scenario B** which are a combination of CQs.

4 EVALUATION

The evaluation of the MindSpaces KG was twofold. On the one hand, we evaluated the consistency and completeness of the MindSpaces KG; we did this with two different evaluation methods. Firstly, we evaluated the completeness of the MindSpaces KG, by defining a set of CQs that the KG must be able to answer with the information it contains (subsection 4.1). Secondly, we evaluated the consistency of the MindSpaces KG by testing if it obeys a set of SHACL constrains (subsection 4.1). On the other hand, the evaluation of the reasoning mechanism was performed by computing the precision-recall-F1 scores used for reasoning systems (subsection 4.2).

4.1 Competence and Consistency of the Knowledge Graph

The completeness of the MindSpaces KG was evaluated through a set of CQs assembled during the formation of the official ontology requirements specification document (ORSD) (Suárez-Figueroa et al., 2009). For this reason, before constructing the MindSpaces KG, we asked from users to define a set of questions that they would like from the MindSpaces KG to contain as knowledge and be able to answer. The users were architects from the School of Architects of the Aristotle University of Thessaloniki[4], either undergraduate-master students or professors, and architects or designers from Zaha Hadid Architects[5]. In total a number of 83 CQs was collected, the complete list of CQs can be found here[2].

Based on the fact that we constructed the MindSpaces KG on the aforementioned set of CQs, this means that if any of the CQ is translated into a SPARQL counterpart, our KG would answer the question with the information it contains. For this reason, we translated each CQ into a SPARQL counterpart and we expected to return the desired information. The completeness of the MindSpaces KG was found adequate, as each CQ when translated into

[3]http://160.40.52.169:6161

[4]https://www.auth.gr/school/arch/

[5]https://www.zaha-hadid.com/

a SPARQL counterpart returned the desired information.

Additionally to the CQs, we performed a validation procedure in order to inspect the syntactic and structural quality of the metadata in the KB and to check the consistency of them. The consistency of the MindSpaces KG was found adequate, as out of 12 SHACL rules, from which 4 referred to object type properties and 8 to data type properties, none of them returned any invalidation of the rule. Moreover, we checked if instances exist which belong to intersection of classes, as we did not desire such a case, and there were not any.

4.2 Knowledge Retrieval Metrics

The evaluation of the reasoning mechanism was conducted using the precision, recall and F1-score used for reasoning systems (Equations 1, 2 and 3), over the two real-life scenarios presented in 3.3 (i.e., Scenario A and Scenario B).

$$precision = \frac{|\{RelevantInstance\} \cap \{RetrievedInstance\}|}{|\{RetrievedInstance\}|} \quad (1)$$

$$recall = \frac{|\{RelevantInstance\} \cap \{RetrievedInstance\}|}{|\{RelevantInstance\}|} \quad (2)$$

$$F1 = 2 * \frac{recall * precision}{recall + precision} \quad (3)$$

Retrieved Instances are considered all the images for which the reasoning mechanism, *did not* return an error when we casted a question to retrieve information for them.

Relevant Instances are considered all the images for which the reasoning mechanism, managed to return some information, when we casted a question to retrieve information for them.

We denote by $Retrieved_A$, $Retrieved_B$ the number of retrieved images for Scenario A and Scenario B, respectively. $Relevant_A$, $Relevant_B$ are the numbers of relevant images for Scenario A and Scenario B, respectively. Next, $precision_A$, $precision_B$ are the precision scores for Scenario A and Scenario B, $recall_A$, $recall_B$ are the recall scores for Scenario A and Scenario B, and $F1_A$, $F1_B$ are the F1 scores for Scenario A and Scenario B, respectively.

The dataset on which we evaluated our reasoning mechanism contains a set of 1200 images, and can be found here[2]. All images were considered Retrieved, meaning our reasoning mechanism did not return any error. Thus, $Retrieved_A = Retrieved_B = 1200$. The same does not hold for the relevant images, for both scenarios, as for Scenario A the $Relevant_A$ images were 1157, and for Scenario B the $Relevant_B$ images were 1142.

Based on the aforementioned numbers the precision, recall and F1-scores for both scenarios can be found in Table 1. Notice, the results are rounded to four decimals.

Table 1: Precision, recall and F1-scores for Scenario A and Scenario B.

	Precision	Recall	F1
Scenario A	0.9649	1.0	0.9821
Scenario B	0.9632	1.0	0.9812

5 DISCUSSION & CONCLUSION

In this paper, we presented the MindSpaces KG, a KG for representing emotions-relevant and functional design for indoor and urban adaptive design. The MindSpaces KG is populated with emotional, physiological, visual, and textual measurements, for the development of adapting environments. It is mostly oriented for artists, designers and architects, and its purpose on one hand is for redesigning indoor workspaces for the workers to feel more comfortable, and on the other hand to improve functionality, mobility, attractiveness of spaces by taking into respect the cultural heritage and environment. But other users can also utilize the MindSpaces KG, for instance the ideas of citizens of a city can be very useful in redesigning urban areas. Moreover, the MindSpaces KG offers a reasoning mechanism to access the information in the KG in a convenient way. For this reason, it retrieves helpful information for real-life scenarios. The scenarios that were used to develop the reasoning mechanism, were provided by experts.

The final step was to evaluate the completeness, the consistency and the reasoning mechanism of the MindSpaces KG. The completeness of the KG (Section 4.1) was evaluated with CQs, which were collected by domain experts. Then, we translated each CQ into a SPARQL counterpart, and checked each one's results, proving that our KG is able to provide information to all users. The consistency of the KG (Sections 4.1) was evaluated with a set of custom constraint rules created and enforced upon the KG where no violation warnings were detected.

The high F1 scores achieved both for Scenario A (98.21%) and Scenario B (98.12%), show that it can be used as an individual mechanism for helping users, by providing insightful information. Additionally, for the relevant instances that were missed, both for Scenario A and Scenario B, we can comment that this happened because these two scenarios require an image as an input, which is analyzed and the Top-3 segmented labels (i.e., Top-3 object labels that are most likely contained in the image) are considered in order to find similar images from the KG (see subsection 3.3). Therefore, the missed relevant images contained Top-3 segmented labels which did not exist in any image in our KG simultaneously.

As for future work, our plan is to enrich the MindSpaces KG with domain-knowledge from relevant Semantic Web KGs, such as ConceptNet (Speer et al., 2017) and WordNet (Fellbaum, 2010), and compare the knowledge in the MindSpaces KG with other existing KG related to the indoor and urban adaptive design. Moreover, we will investigate and expand the quantity of real-life scenarios that the reasoning mechanism can support. Finally, we also plan to create a more friendly user interface, as at the moment the queries are formulated through SPARQL.

ACKNOWLEDGEMENTS

This work has been supported by the EC-funded project MindSpaces (H2020-825079)

REFERENCES

Andreadis, S., Moumtzidou, A., Apostolidis, K., Gkountakos, K., Galanopoulos, D., Michail, E., Gialampoukidis, I., Vrochidis, S., Mezaris, V., and Kompatsiaris, I. (2020). Verge in vbs 2020. In *International Conference on Multimedia Modeling*, pages 778–783. Springer.

Asim, M. N., Wasim, M., Khan, M. U. G., Mahmood, N., and Mahmood, W. (2019). The use of ontology in retrieval: a study on textual, multilingual, and multimedia retrieval. *IEEE Access*, 7:21662–21686.

Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web. *Scientific american*, 284(5):34–43.

Chi, N.-W., Jin, Y.-H., and Hsieh, S.-H. (2019). Developing base domain ontology from a reference collection to aid information retrieval. *Automation in Construction*, 100:180–189.

Fellbaum, C. (2010). Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer.

Hyvönen, E. (2012). Publishing and using cultural heritage linked data on the semantic web. *Synthesis lectures on the semantic web: theory and technology*, 2(1):1–159.

Kumar, V. R. S., Khamis, A., Fiorini, S., Carbonera, J. L., Alarcos, A. O., Habib, M., Goncalves, P., Li, H., and Olszewska, J. I. (2019). Ontologies for industry 4.0. *The Knowledge Engineering Review*, 34.

Lopes, D. M. (2007). Shikinen sengu and the ontology of architecture in japan. *The Journal of aesthetics and art criticism*, 65(1):77–84.

Lynch, K. et al. (1960). The image of the city (vol. 11).

Munir, K. and Anjum, M. S. (2018). The use of ontologies for effective knowledge modelling and information retrieval. *Applied Computing and Informatics*, 14(2):116–126.

Pistola, T., Georgakopoulou, N., Shvets, A., Chatzistavros, K., Xefteris, V.-R., García, A. T., Koulalis, I., Diplaris, S., Wanner, L., Vrochidis, S., et al. (2022). Imageability-based multi-modal analysis of urban environments for architects and artists. In *International Conference on Image Analysis and Processing*, pages 198–209. Springer.

Qiu, W., Li, W., Liu, X., and Huang, X. (2021). Subjective street scene perceptions for shanghai with large-scale application of computer vision and machine learning. Technical report, EasyChair.

Raven, D., de Boer, V., Esmeijer, E., and Oomen, J. (2020). Modeling ontologies for individual artists. *Vrije Universiteit Amsterdam*.

Schneider, A. (2020). Alternatives: World ontologies and dialogues between contemporary arts and anthropologies. In *Alternative Art and Anthropology*, pages 1–26. Routledge.

Speer, R., Chin, J., and Havasi, C. (2017). Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.

Suárez-Figueroa, M. C., Gómez-Pérez, A., and Villazón-Terrazas, B. (2009). How to write and use the ontology requirements specification document. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, pages 966–982. Springer.

Vassiliades, A., Bassiliades, N., Gouidis, F., and Patkos, T. (2020). A knowledge retrieval framework for household objects and actions with external knowledge. In *International Conference on Semantic Systems*, pages 36–52. Springer, Cham.

Wagner, A. and Rüppel, U. (2019). Bpo: The building product ontology for assembled products. In *Proceedings of the 7th Linked Data in Architecture and Construction workshop (LDAC 2019)', Lisbon, Portugal*, page 12.

Yu, B. (2019). Research on information retrieval model based on ontology. *EURASIP Journal on Wireless Communications and Networking*, 2019(1):1–8.