Mini-Map: Mixed-Initiative Mind-Mapping via Contextual Query Expansion

Ting-Ju Chen*, Sai Ganesh Subramanian†, and Vinayak R. Krishnamurthy‡
Texas A&M University, College Station, Texas, 77843

We present Mini-Map, a novel web-based workflow for mixed-initiative mind-mapping using query expansions enabled by ConceptNet — a large graph-based representation of "commonsense" knowledge. Mind-maps are now widely recognized as thinking tools for creative tasks such as conceptual design. Nonetheless, creating a mind-map can be challenging especially for beginners either due to low domain knowledge or personal inhibition. Despite extensive knowledge databases on the web, there is currently limited computational support to augment users' ability to fully explore the vast information sources available at their disposal. Our work demonstrates that by using the rich relational ontology offered by ConceptNet, we can faithfully emulate human-like collaborative behavior toward useful exploration of ideas in early design. We present a comparative analysis of human-human collaborative mind-mapping and the Mini-Map workflow.

I. Introduction

This paper seeks to enable and understand how humans and computers can creatively *co-generate* ideas with equal participation to explore concepts around a design problem. To this end, we present Mini-Map, a digital workflow for **m**ixed-**ini**tiative mind-**map**ping wherein a human designer and a computer program take turns to create a mind-map for a given design problem. The workflow is designed to gamify the process of mind-mapping where a designer is collaboratively playing a game with an Intelligent Agent (AI). The AI-player is powered by ConceptNet [1–4], which is presently one of the largest commonsense knowledge base that covers assertions between concepts through rich relational ontology.

Mind-mapping is widely used for early-stage design brainstorming by creating a network (mostly tree-like structures) of concepts or ideas surrounding a central idea. From this viewpoint, mind-maps are primarily tools for *problem understanding* before problem solving in that they allow an unconstrained exploration of a variety of ideas along with the relationships between those ideas in a hierarchical fashion. The rules of creating a mind-map are rather simple: one starts with a central idea and creates two to three *branches* leading to related ideas repeating the process for each newly added idea. Furthermore, there are now several digital tools [5–12] that can be used to create and document mind-maps. Despite this, creating such maps can be challenging, especially for novices for a number of reasons including but not limited to (a) lack of knowledge in the domain of a given design task, (b) unstructured and ill-defined task that requires self-creativity and adaption, (c) inability to readily recall known concept, (d) inability to envisage relationships across concepts due to complex couplings between variables, and (e) high dimensionality of the problem space.

Our work draws from the notions of mixed-initiative co-creativity elaborated by Yannakakis et al. [13] and study a simple yet powerful workflow in the context of *exploratory* creativity in early design. Specifically, we present the following ideas. First, we introduce a new work-flow — Mini-Map — that re-formulates mind-mapping as a collaborative game between human and computer collaborators. Second, we present a novel algorithm that is powered by the relational ontology offered by ConceptNet [1] in conjunction with statistical, topological, and temporal rules for adding concepts as nodes to the mind-map. Finally, we conduct a human subject study where we compare human-human and human-computer collaboration for mind-mapping. Our evaluation is based on the adaptations of well-established creativity metrics proposed in design research [14]. This comparison demonstrates the efficacy of our approach for digital mind-mapping and offers some critical insights into how digital mind-mapping can be advanced through mixed-initiative approaches.

*Graduate Research Assistant, J. Mike Walker '66 Department of Mechanical Engineering, 3123 TAMU, College Station, TX 77843-3123.
†Graduate Research Assistant, J. Mike Walker '66 Department of Mechanical Engineering, 3123 TAMU, College Station, TX 77843-3123.
‡Assistant Professor, J. Mike Walker '66 Department of Mechanical Engineering, 3123 TAMU, College Station, TX 77843-3123.

A. Contributions

We make three primary contributions in this paper. (a) Our work is first of a kind to present human-computer collaboration for an unstructured task scenario. (b) By performing multiple iterations of the proposed algorithm, we present knowledge on how computer can perform better in a mixed-initiative task. (c) We conduct a user evaluation comparing Human-Human (**HH**) and Human-Computer (**HC**) collaboration using computational and subjective methods.

II. Related Works

A. Mind-maps for Ideation

Mind-maps can enable effective thinking and learning [15, 16] through rapid idea expression and divergent exploration [17]. Their general structure makes them useful for a wide variety of applications ranging from document drafting [18], project planning [19], and decision making [20]. Zampetakis et al. [21] discuss the utility of mind-maps in the learning process of engineering students. Specifically during design ideation, they can be immensely useful for reflection, note-taking, idea communication, and idea synthesis while reducing the cognitive load accompanied with retrieving and maintaining diverse knowledge elements [22, 23]. Apart from their structure, mind-maps can be further enriched by colors, images, arrows, and dimension to reflect personal interest and individuality [24]. Given our specific focus mixed-initiative interactions for mind-mapping, we have constrained the variations in the visual variables, restricting the idea representation to a a simple node-link diagram with words (and potentially phrases) as the node content.

B. Digital Mind-mapping

Several works have investigated the effect of interactivity offered by digital tools on idea generation and collaboration in mind-mapping. Buisine et al. [25] showed that there is no significant difference in idea production between digital tabletop-based and traditional paper-and-pencil mind-mapping workflows. However, they also report better collaborative dimensions, especially by leading to better-balanced contributions from the group members. Faste and Lin [26] evaluated numerous existing mind-mapping software applications, performed ethnographic studies with a variety of users, and developed a framework of principles to guide future development of digital mind-maps. Orehovavcki et al. [27] address the quality in use of Web 2.0 applications for mind-mapping purposes. They found that participants feel highly satisfied with Web 2.0 applications that have the following attributes: ease of use, effectiveness, controllability, interactivity, navigability, customizability, efficiency, information content coverage, understandability, and reliability.

C. Automated Mind-mapping

Several works [28–31] have investigated text mining and machine learning approaches for mind-map generation. However, these are fundamentally different from our approach as they focus on completely automated generation of maps. To the best of our knowledge, ours is the first approach that demonstrates a successful *collaborative behavior* between human and computer agents for mind-mapping.

D. Identification of Gaps

There have been several efforts [32, 33] in enabling mixed-initiative systems in the domain of *game design*. Most of these works are essentially targeted toward procedural content generation in computer games. While there are recent advances in aerospace design domains such as the *Daphne* system [34], there is currently little known about how human-computer co-creativity could be materialized for engineering and product design domains in early conceptualization. Our work is a step toward filling this knowledge gap.

III. Technical Approach

Our technical approach is driven by the observation that in a digital setting where the user has access to vast knowledge databases, the cognitive processes underlying mind-mapping could be effectively supported by systems that allow for not just querying singular concepts from the database but by allowing the user to expand upon those queries. Second, mind-maps lend themselves to a natural graph based structure with concepts (words/phrases) as nodes and their relations as the links or edges. Based on this, our approach is to use ConceptNet, a semantic network that contains a graph-based representation with nodes representing real-word concepts as natural language phrases (e.g. *block a door*, *be abandoned*, etc.), and edges representing semantic relationships.

A. Problem Formulation

We assume a mind-map to be a tree (an acyclic graph), say $M(V, E)$ with a set of nodes V_M and edges E_M and the construction of a mind-map can be viewed as an iterative sequence of breadth-first (i.e. exploring the aspects around a problem) and depth-first (i.e. adding detailed and concrete ideas). The root node of this tree is the central problem around which and we aim to emulate this behavior in a collaborative fashion wherein a human and an automated agent are collectively working to diversify ideas around the root node. We formulate process of mixed-initiative mind-mapping as a collaborative interactive game that is being played by a human and an intelligent agent (AI). Given a central topic, the aim for the human and AI collaborators is to expand the map by exploring different aspects of the topic (breadth-first), and refine each of these aspects by adding detailed concepts (depth-first). In order to do this, the collaborators are given a seed topic as the root node of M and the human and AI players take turns to add nodes to M. Furthermore, each player is allowed exactly one node addition per turn.

Given any state of M (the simplest being a given root node representing the central idea), there are two primary algorithmic steps that are needed for creating a mini-map. The first step is *target search* wherein the AI needs to determine a target node (a node where a new node will be added). Subsequently, the second step is *content generation* wherein the AI needs to define the content (words and phrases) that will be added to the target node to ultimately expand the space of ideas around the central problem.

B. Approach

1. Target Search

One of the most important aspects of mind-mapping that we wish to capture is divergent idea exploration that is key in ideation processes. Additionally, we also wish to support detailed exploration of a given chain of ideas (formally this would be a path from the root node to a given leaf node in M). For this, our main goal is to model the behavior of the AI player in such a way that M evolves in a breadth-first fashion in the initial stages of mind-mapping and transitions into a depth-first target selection during the later stages of the evolution of M. In this work, we experimented with several alternatives for target search including completely random target search, search based on the topological and temporal evalution of M, and search based on the level of expansion on a node-by-node basis. These are described in detail in subsequent sections.

2. Content Generation

Drawing from the notion of query expansion and pseudo-relevance feedback, we develop the method for content generation via ConceptNet [35], a large graph of concepts where two concepts are directly connected with an edge if there is a semantic relation between them. Our choice is based the the potential of ConceptNet for making more complex, multi-step inferences that are helpful in query formulation and the identification of poorly performing queries. Other existing methods for query expansion rely mostly on content-retrieval analysis (Wiki) or lexical-semantic database such as WordNet [36]. However, although WordNet is widely used to provide hypernym relations in a hierarchical fashion, ConceptNet allows for the organization of related concepts based on a diverse set of relations (such as "*IsA*", "*HasA*", "*UsedFor*", "*CapableOf*") resulting in a broader scope of queries that WordNet synsets would not allow. Using this feature of ConceptNet, we developed several alternatives for selecting the content for node addition such as global median, relation-dependent median, and context-dependent weighting schemes. These are detailed in the subsequent sections.

C. Algorithmic Iterations

In this work, we take an iterative approach for the development of the Mini-Map algorithm. Specifically, our final algorithm for Mini-Map was developed based on three prior iterations. We combined different alternatives for target search and content generation in each iteration and conducted small scale pilot studies with users. The observation of the resulting mind-maps and issues raised by users allowed us to refine each subsequent iteration culminating in our final algorithm. Below, we provide a detailed account of our iterations and the key insights we gained in terms of the AI-behavior we observed.

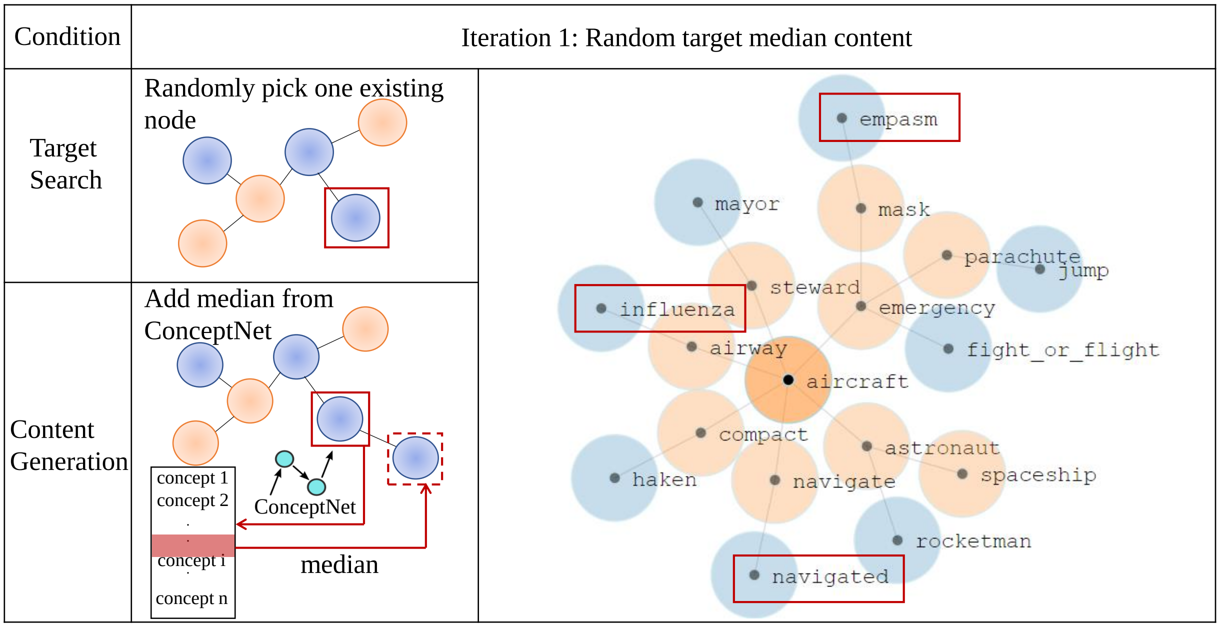

Fig. 1 Illustration of iteration 1 in our approach. The resulting mind-map shows several disadvantages: users hard to keep up responses from AI because of random target search; the AI is unable to generate multiple ideas to one concept; and the generated content might be out of context (which are highlighted in red boxes: *influenza*, *navigated*, *empasm*, etc.)

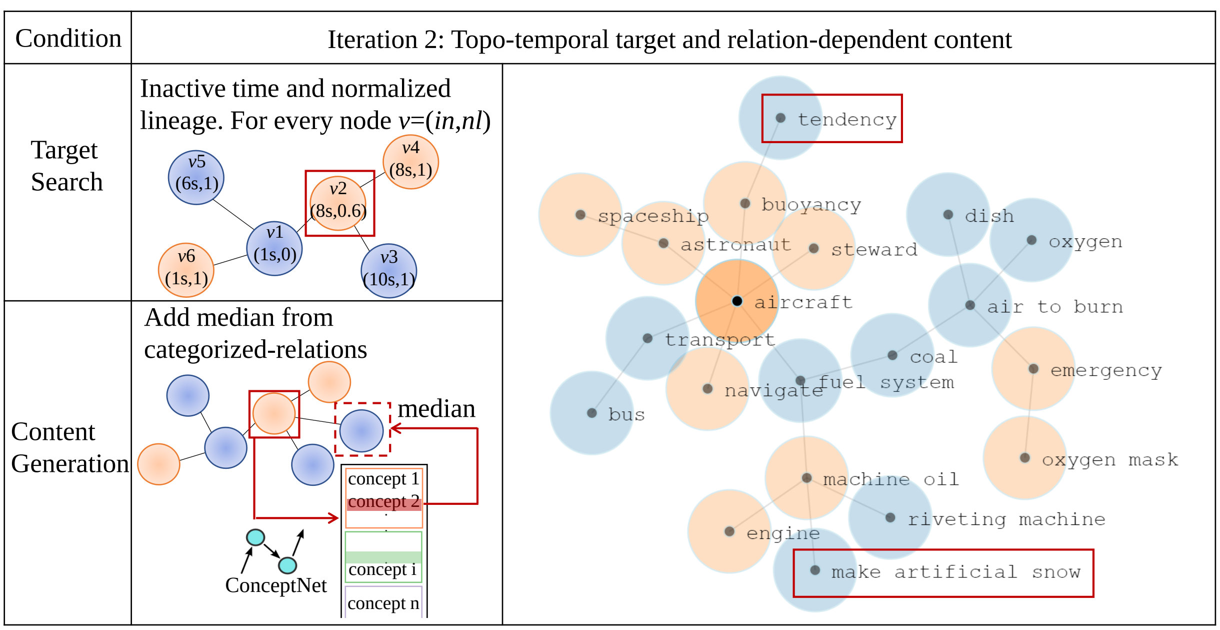

Fig. 2 Illustration of iteration 2 in our approach. The resulting mind-map shows a main disadvantage: the AI generates content with little or no relevance to the context of current mind-map (which are highlighted in red boxes: *buoyancy* is a type of *tendency*, *make artificial snow*, etc.)

1. Iteration 1: Random Target and Median Content

In our first iteration, we applied a simple rule for selecting targets to add a new node to (Figure 1). Given the current state of a map M with nodes V_M and edges E_M, we first randomly select a concept $v \in V_M$. Then, we

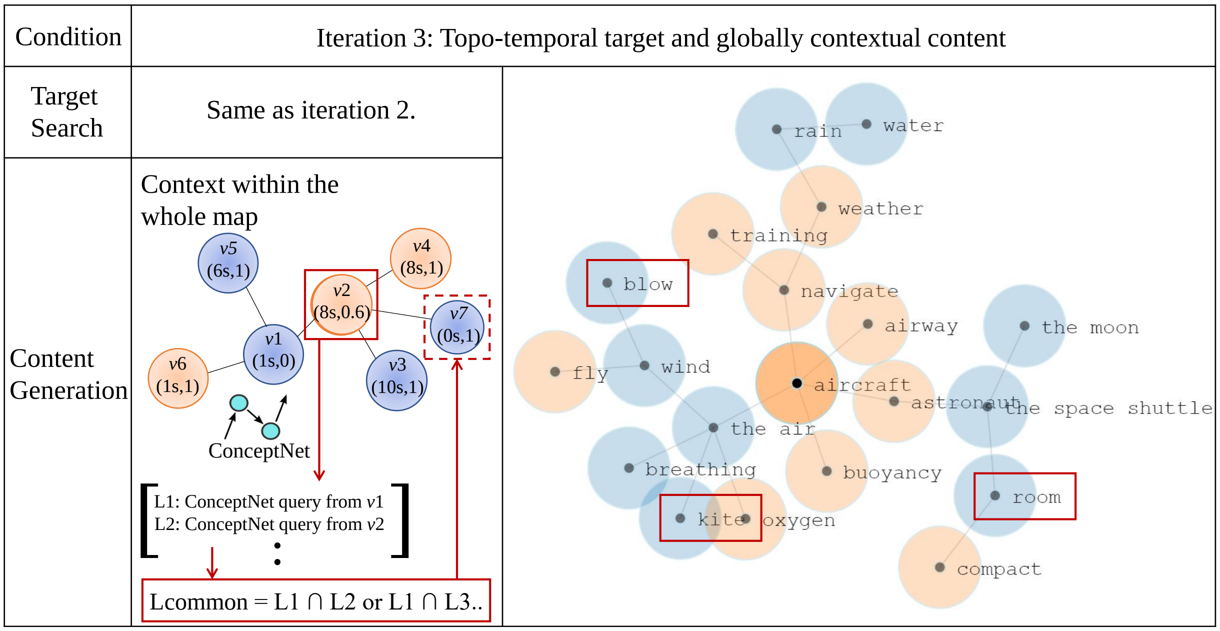

Fig. 3 Illustration of iteration 3 in our approach. The resulting mind-map reveals several disadvantages: the AI can not response to last added node due to measure from inactivity in target search, and the generated contents become abstract as mind-map grows (which are highlighted in red boxes: *kite* from *the air* and *fly*, *blow* from *wind* and *breathing*, *room* from *wind* and *space*).

retrieve a set $Q = \{q_i | (v, q_i) \in E(C), i \in [1, k]\}$ and sort it in decreasing order of the edge weights in ConceptNet (i.e. $w(v, q_1) \geq w(v, q_2) \geq \cdots w(v, q_k)$). For content generation, we select the concept with the median edge weight and add to the selected node. The computation terminates when there are no further concepts where the AI can add a node (this is guaranteed since there is only one median element in a given set). Given a target node, the seemingly obvious choice would be to query the target node and generate content based on the highest weighted results given by ConceptNet. This approach clearly leads to more abstract concepts being added to the map as it grows — a scenario that is in contrast with how mind-mapping is typically performed. For instance, for a central topic such as "*Aircraft*", adding a node "*fly*" leads the intelligent agent to add the word "*sky*" followed by "*blue*" and "*a colour of the rainbow*)". The primary reason for choosing the median weighted edge is the observation that the edge weight distribution of the linked concepts in ConceptNet is skewed. Hence, taking the median instead of other statistical aggregates such as mean results in a more robust result in terms of preserving the relatedness as well as diversity of the added concepts.

Our naive approach provided some interesting results (Figure 1) for the central topic *Aircraft*. As can be seen, the method is able to produce reasonable maps that can potentially lead to technically feasible solutions to the problems. However, we also observe that the AI partner sometimes adds nodes that may not be relevant to the current design context (e.g. "*navigated*" in response to "*navigate*") or related concepts that may not be useful in arriving at solutions (e.g. "*fight-or-flight*" in response to "*emergency*"). The second limitation of the naive method is the fact that the AI partner can not add more than one node to a given existing concept in the existing map. This is merely the result of the fact that there is only one median that we can choose for a given ConceptNet query

2. Iteration 2: Topo-Temporal Target and Relation-Dependent Content

In order to overcome the obvious shortcomings of random target search, our second iteration investigated a method for target search based on the topological and temporal evolution of a given mind-map. Here, for each node $v \in V_M$, we introduce two penalty terms, namely, inactivity ($in(v)$) and normalized lineage ($nl(v)$). Here, $in(v)$ is computed based on the time elapsed since last activity that occurred at node v (i.e. addition of a child node). Upon addition of a child node to v, we reset $in(v) = 1$ and linearly decrease $in(v)$ at every computation cycle until it reaches 0. Further, $nl(v)$ is computed by normalizing total count of all descendents of v. It is 1 for a leaf node of M and 0 for the root node. By applying a threshold to these penalties, the AI determines the target nodes that lie below the threshold.

In order to enable the AI to add multiple nodes to target node, we further subdivide a given query list into the defined relation-categories and add median-weighted concepts from each of these categories starting from the highest to lowest weights. This also allows the AI to preserve both the relevance (higher weights) and novelty (lower weights) in the resulting mind-map (Figure 2). However, we found that the content generation with this approach usually led to addition of concepts with little or even no relevance to the context of current mind-map. For instance, "*buoyancy*" leads to "*tendency*" and "*machine oil*" leads to "*make artificial snow*".

3. Iteration 3: Topo-Temporal Target and Globally Contextual Content

Considering the behavior results from iteration 2, we modified our content generation rules keeping the target search the same as iteration 2. Our intent was to take into account the context of the currently evolving mind-map during the generation of new content. For this, at any given state of the mind-map, we compile a list of queries that are common across all concepts in the mind-map. Subsequently, we use the subdivided median weight approach similar to iteration 2 for this common list. This ensures that the generated contents share common properties from the existing context in a mind-map (Figure 3). From iterations 2 and 3, the measure of time in deciding potential targets make the AI not response to any last added node by user, since they just got active. This phenomenon results in user feedback such as *"make me lose my train of thought"*. Also, we found out if choosing from overlapped concepts from the whole context, the results are also going to become abstract as the mind-map grows, hence lose focus. For example, *kite* was generated based on *fly* and *the air*, *blow* from *wind* and *breathing* (Figure 3).

D. Final Mini-Map Algorithm

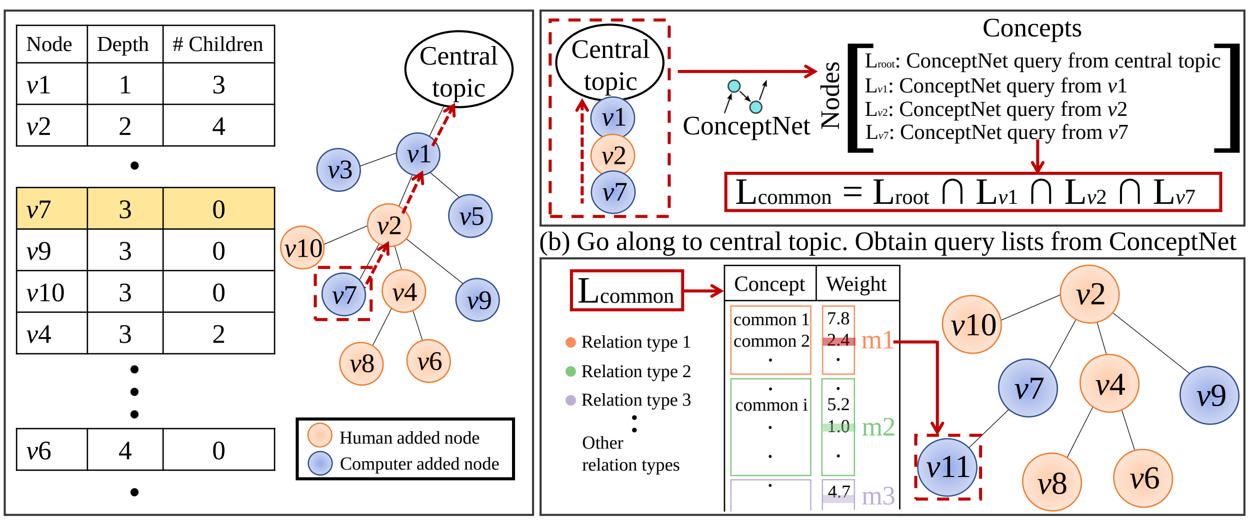

(a) Mini-Map target search (c) Categorize & sort. Add content from median weight

Fig. 4 Illustration of the algorithm to generate a new node from computer using retrieved weighted relations through ConceptNet Open Data API. (a) Considering a branch from current mind-map, find "where" to add based on sorted *depth* and *number of children* of all nodes. (b) Trace back to the central topic (root). Retrieve query results (L) from ConceptNet with respect to each node. Then find the common concepts across the retrieved results (L). (c) Categorize the common concepts and sort it by weights in ConceptNet. Computer adds a new node from median weight.

1. Target Search: Expansion Threshold

The main issues with the previous iterations of the target search was the imbalanced addition of nodes to the map leading to heavily biased exploration of ideas. The temporal approach for node addition further pronounced this imbalanced map evolution leading to longer chains (meaning more detail-based exploration). In our final iteration, we developed a target search strategy based on *expansion threshold*. The basic idea is to enforce a breadth-first strategy (i.e. nodes closer to the root get preference) in a manner that ensures a bare minimum expansion of a given node.

Specifically, the AI would traverse the mind-map starting from the root and will prefer the target closest to root with less than an *expansion threshold* (i.e. a pre-determined number of child nodes). Based on experiments, the expansion threshold of three child nodes was determined to be reasonable. (Figure 4(a), Algorithm 1). Note that this algorithm has a guaranteed termination criterion since there are always leaf nodes in M (i.e. nodes with no child nodes).

Algorithm 1 Target Search

Input: Current mind-map $M = (V_M, E_M)$
Input: Bi-ordered Index Set $X \in \mathbb{Z}^3$
Output: Target Node $v_T \in V_M$
Output: Updated Bi-ordered Index Set $X \in \mathbb{Z}^3$

if $X = \emptyset$ **then**
  $X \leftarrow \emptyset \in \mathbb{Z}^3$
  for $v_i \in V_M$ **do**
    $x_1 \leftarrow$ depth of v_i
    $x_2 \leftarrow$ number of children of v_i
    $x_3 \leftarrow i$
    $\mathbf{x} \leftarrow (x_1, x_2, x_3)$
    $X \leftarrow X \cup \{\mathbf{x}\}$
  end for
  $X \leftarrow Sort(X,\ descending\ order\ in\ x_1)$
  $D_{max} \leftarrow$ maximum depth of M
  for $d \in [1, D_{max}]$ **do**
    $X \supset X_d \leftarrow \{\mathbf{x} = (x_1, x_2, x_3) \in X | x_1 = d\}$
    $X_d \leftarrow Sort(X_d,\ descending\ order\ in\ x_2)$
  end for
end if
for $d \in [1, D_{max}]$ **do**
  $X \supset X_d \leftarrow \{\mathbf{x} = (x_1, x_2, x_3) \in X | x_1 = d\}$
  for $\mathbf{x}^j = (x_1^j = d, x_2^j, x_3^j) \in X_d$ **do**
    if $x_2^i < 3$ **then**
      $i \leftarrow x_3^j$
      $v_T \leftarrow v_i \in V_M$
      $x_1^i \leftarrow x_1^i + 1$
      return v_T, X
    end if
  end for
end for

2. Content Generation: Path Dependent Context

In iteration 3, the primary disadvantage of using the whole mind-map to determine context-dependent content is that the resulting generated content is abstract. This is natural since the concepts in the whole mind-map, since reasonably dissimilar, are bound to be connected only be the most abstract concepts. For instance, *kite* was generated based on *the air* and *wind* (Figure 3). While we did aim at contextual relevance, we also wanted to maintain an appropriate level of detail as the map evolves. To achieve this, we implemented our final content generation algorithm by defining context based on the pat rather than the whole map. Specifically, given a target node $v_T \in V_M$, the idea is to create a doubly sorted list, L_{common}, of all concepts that are common to all concepts along the shortest path from v_T to the central idea. Here, by doubly-ordered, we mean that: (a) the all entries in L_{common} are listed categorically based on the 25 relationship types provided by ConceptNet, (b) the entries for each relationship type are further ordered with decreasing weight of the relationship, and (c) the relations are themselves ordered from highest to lowest total weights ((Figure 4(b), steps 1 to 16 in Algorithm 2)).

3. Content Generation: Relation Nodes and Median-Weighted Nodes

Similar to the strategy in iteration 2, we retained the capability of the AI to add multiple nodes from within the same relationship. Given a list of common ConceptNet queries along a path, we achieve this by subdividing the list into relation-categories for allowing AI to be able to add multiple nodes to an existing concept and preserving both the relevant (higher weights) and novelty (lower weights) (Figure 4(c)). In a typical mind-map, the nodes connected to the central idea usually determine the type of details that will be consequently added. For instance, for a topic such as *pollution*, one may first add nodes such as *causes*, *effects*, *mitigation* etc. instead of mentioned what the causes are or what the effects will be or how to mitigate pollution. This is difficult to achieve by simply adding median-weighted queries from ConceptNet. Here, we take a crucial step here; we observed that the relationship types in ConceptNet are themselves helpful to users in organizing their ideas, particularly during early phases of mind-map evolution. We capture this by enforcing the AI to add ConceptNet *relations* as content nodes as the first several main branches from central topic contained in the root node (steps 17 to 22 in Algorithm 2).

Algorithm 2 Content Generation

Input: Target Node $v_T \in V_M$
Input: Relation Descriptor Set $R(v_T)$
Output: Concept $c \in V_C$ from ConceptNet graph $C = (V_C, E_C)$
Input: Relation Descriptor Set $R(v_T)$
1: $P_v \leftarrow$ list of nodes on the shortest path between v_T and Root Node
2: $L_{common} \leftarrow \emptyset$
3: **for** $p_i \in P_v$ **do**
4: $\quad L_i \leftarrow \{(c_j, r_j, w_j)|(p, c_j) \in E_C\}$ (r_j is the type of relation, w_j is edge weight))
5: **end for**
6: $L_{temp} \leftarrow \bigcap_{i=1}^{|P_v|} L_i$
7: $L_{temp} \leftarrow Sort(L_{temp}, alphabetical\ order\ of\ r_i)$
8: **for all** $r \in R$ **do**
9: $\quad L_{temp} \supset L_r \leftarrow \{l_i = (c_i, r_i, w_i) \in L_{temp}|r_i = r$
10: $\quad L_r \leftarrow Sort(L_r, descending\ order\ of\ w_i)$
11: $\quad W_r \leftarrow \sum w_i \forall l_i \in L_r$
12: **end for**
13: $RL \leftarrow \{L_{r1}, L_{r2}, \ldots\}$
14: $RW \leftarrow \{W_{r1}, W_{r1}\}$
15: $L_{common} \leftarrow Sort(RL, descending\ order\ of\ RW)$
16: $R(v_T) \leftarrow Sort(R(v_T), descending\ order\ of\ RW)$
17: **if** Depth of $v_T = 0$ **then**
18: $\quad c \leftarrow r_1 \in R(v_T)$
19: **else**
20: $\quad c \leftarrow c_j \in L_{r1}|w_j = meadian(w_j \forall l_j \in Lr1)$
21: $\quad R(v_T) \leftarrow R(v_T) - \{r_1\}$
22: **end if**
23: return c, $R(v_T)$

IV. Implementation Details

A. Front-end Design

1. Visual Encoding

Similar to traditional mind-maps, the central topics are always reasonably sized. From main branches to details, we encode varying font size and color gradient to visually represent the emphasis of the information. Ideas added from different users would be given different color schemes, which helped users to recognize the thought flow easily. Ideas are spatially organized in the mind-map using forced-links structure in D3JS. Such force-directed layout can help produce elegant spreading of nodes and reasonable visibility of links even with large dataset.

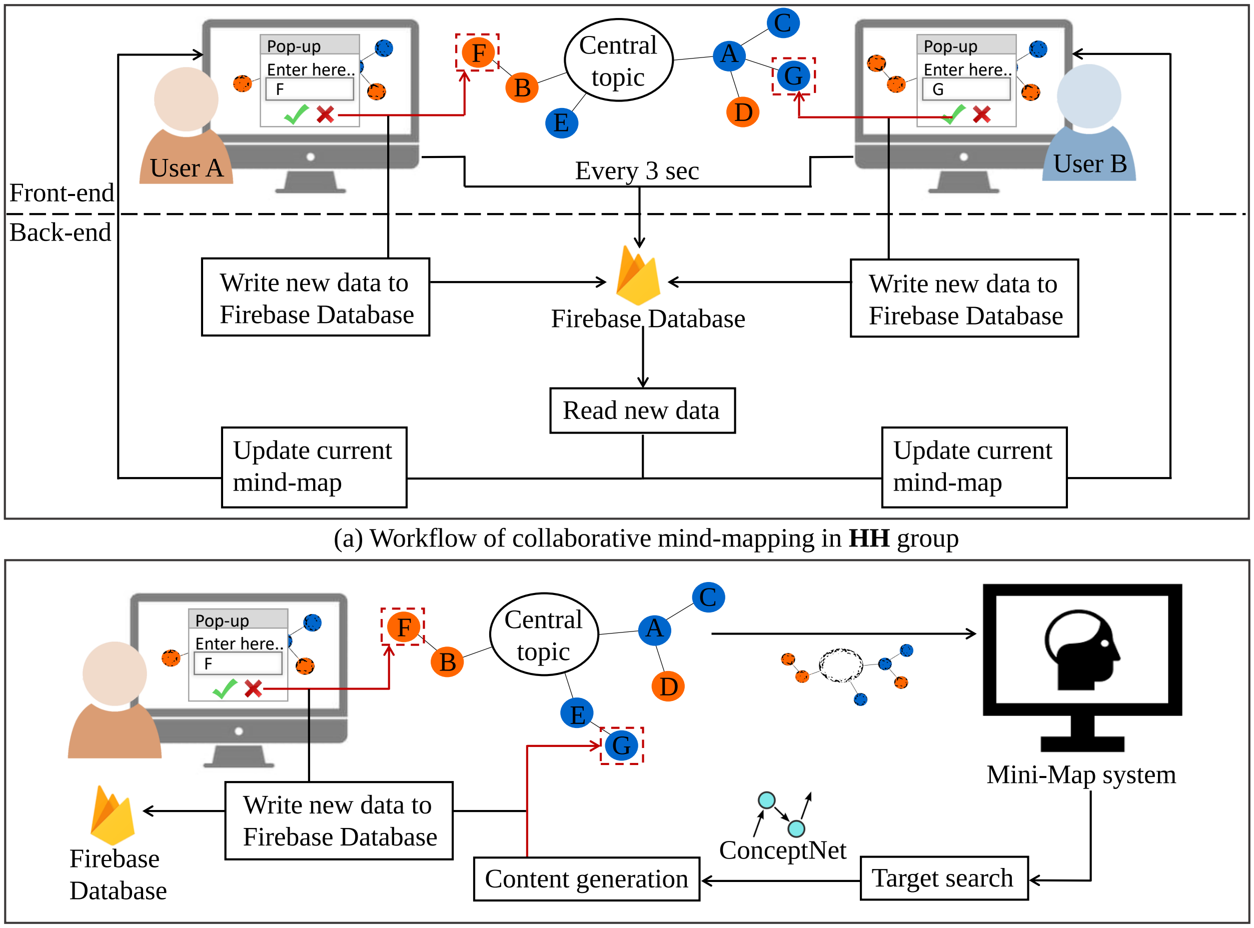

(a) Workflow of collaborative mind-mapping in **HH** group

(b) Workflow of Mini-Map in **HC** group

Fig. 5 Flow showing the front-end (user interaction) and back-end (data processing) pipeline in HH and HC.

2. Interactions

We designed a game-like interaction for users to collaboratively create a mind-map; users take turns to add nodes. In each turn, users are allowed to add only one node by double-clicking on any of the existing nodes. An input dialog pops up after double-clicking in the editor workspace. Once a valid answer submitted, our system creates a new child node with the double-clicked one being its parent node. A link is created automatically between them. This offered users minimal manipulation in the construction of a tree type structure. On the other hand, while it is not your turn, an layer appears on top of the interface to prevent any interactions.

B. Back-end

1. MiniMap.js Library

Mini-Map used a JSON data structure specifically designed for mind-map, which included two major components: node object and link object. Node object has attributes including position data (x,y), properties (unique label, parent, children, depth, etc.) and time series data (added time, modified time). Link object has attributes including source node, target node and a unique label. Whenever a mind-map is imported or loaded into the workspace, the JSON file is read by Mini-Map and corresponding mind-map regenerated.

2. ConceptNet API

We incorporated JSON-LD API (Linked Data structure) offered by ConceptNet Open Data API in Mini-Map. The API allows for querying a word or a phrase through HTTP request and provides a URL to a page containing all related words and phrases to the query in a linked data format. For each concept in the linked data, */en/* stands for its language code, *rel* contains corresponding relation such as "*UsedFor*", "*CapableOf*", etc., and the strength with which this relation expresses this concept is stored in *weight*. The linked data also provides human-readable label in *start* and *end*.

3. Firebase Pipeline

Mini-Map was implemented using Firebase Database REST API written in javascript (Figure 5). Each mind-map data is stored in Firebase Realtime Database with corresponding user ID. In a collaborative setting where two users are working with the same mind-map, Mini-Map employs a listener function to scan Firebase Realtime Database at intervals of 3 seconds. It checks for changes in the mind-map data, and updates the current mind-map based on the latest data found from Firebase.

V. User Evaluation

We conducted a comparative study to understand how humans interact with their collaborator and explored human behavior in a controlled experimental setup. Our study tasks were designed with two major goals in mind. First, we want to gain insight on how our approach compares with a typical human-human collaboration scenario in terms of quality, variety, and novelty of ideas generated through mind-mapping. Second, we want to understand if our approach can truly facilitate a human-like collaboration experience for design applications.

A. Procedure and Participants

With these goals in view, we performed a user study with two distinct groups of participants. The first group (**HH**) comprised of 28 human participants who were asked to create mind-maps in pairs. The second group (**HC**) of participants comprised of 14 individuals who co-created their mind-maps with Mini-Map as the collaborator. Both the **HC** and **HH** participants did **not** know the type of agent they were collaborating with. In order to simulate human-like behavior, we delayed the response from the computer by 5 to 10 seconds in **HC** study.

Our participants were primarily undergraduate and graduate students from engineering, sciences, and architecture majors to help us sample from a wide variety of disciplines and age groups. The total time taken during the experiment varied typically vary between 30 to 35 minutes. After conducting the demographic survey, describing the purpose of the study, basic features of the system were explained and participants were allowed some time to get acquainted with the setup. In addition, participants were encouraged to ask questions for any further clarifications regarding the study. Following this, the participants created mind-maps for given topics. Each group of participants was asked to create only one mind-map for a given topic.

After preliminary study with various topics, we choose the topics *Solar Energy* and *Space Travel* for the user study. *Solar Energy* is a specific topic and we want to test whether our algorithm is able to add constructive nodes in any given context. *Space Travel*, on the other hand, is a general topic which stimulates the user to imagine, perceive and explore various directions. Our goal is to observe the capability of our algorithm to help the user explore and organize their thoughts both in a specific and open-ended domain. The participants were given 10 minutes to work on each of the aforementioned topics. Also, their screen was recorded over the duration of a given mind-mapping task based on their consent.

B. Data Collected

On completion, each participant were asked to complete a feedback questionnaire based on their experience during the mind-map creation tasks. Here, we wanted to understand the differences in perception of the mind-mapping process from our participants in the **HH** and **HC** groups. For this, we conducted a "*peer evaluation*" from each participant at the end of the mind-mapping session. Our peer evaluation was based on the recent work by Gilon et al. [37]. Also, the whole process of mind-mapping was screen recorded with the consent of the participants.

In addition to specific questions about the interface and collaborator, the users were also asked to predict whether their collaborator was a human or computer. The reasons behind every participant's choice were also elicited and collected. The aim of this question was not to mislead the participant in thinking of the collaborator as a human or computer, but to gain insights on the behavioural aspects and how the users perceived their collaborator's typical behavior. This was inquired both to understand the specific human-like attributes that was reflected in the Mini-Map and to develop understanding of particular shortcomings of **HH** interface that could be overcome by the Mini-Map.

C. Metrics

1. Computational Metrics

In order to perform automatic semantic analysis within the two user groups, we used *ConceptNet Numberbatch* [2], a word embedding model that has been shown to be measurably better than word2vec [38] and Glove [39]. After the model had been trained, we can visualize the learned embeddings using t-SNE (t-Distributed Stochastic Neighbor Embedding) [40] for dimensionality reduction (Figure 8). As the vector space was modeled and each word had its unique vector representation, we used Euclidean distance to measure semantic distance within central topic and ideas in the resulting mind-maps, and plotted the frequencies to identify the diversification of ideas.

2. Subjective Metrics

The mind-maps created at the end of the user study were evaluated by two expert raters. The raters selected were unaware of the study design and tasks, and were not furnished with any information related to the general study hypotheses other than final mind-maps created. Also, the mind-maps generated by **HC** and **HH** study were randomized and de-identified before they were shared for further evaluation. Both the raters were senior graduate research students (and potentially faculty members) in engineering and product design disciplines and were asked to rate each mind-map based on well-established metrics. From the mind-map assessment rubric [41], we adapted the metrics Structure, Exploratory, Communication and Extent of Coverage for a comprehensive assessment of the mind-maps. The raters were then asked to evaluate every mind-map based on these metrics on a scale of 1 to 4 (Figure 6).

	Level 1	Level 2	Level 3	Level 4
Structure	Few ideas radiate from center. Not very clear	Some ideas radiate out from center but are not suitable to topic	Ideas radiate out from center in a clear picture that involves imagination and creativity	Ideas provide a complete picture with a high degree of imagination and creativity
Exploratory	Ideas are not connected from most complex to simplest	Some ideas move from most complex to simplest	Ideas are arranged in order of importance from most complex to simplest	Clear and highly effective indication of connection between ideas and center topic
Communication	Limited use of keywords	Key words are used. Average understanding of topic	Good use of key words connected to central topic. Good understanding of topic	Highly effective use of key words. Deep understanding of topic
Extent of Coverage	Limited or ineffective effort to connect main ideas together	Good or adequate effort to connect main ideas together	Effective effort to connect main ideas together	Highly effective effort to connect main ideas together

Fig. 6 Table explaining the metrics Structure, Exploratory, Communication and Extent of Coverage that were given to the inter-raters as a reference for evaluating mind-maps. This was adopted from the mind-map assessment rubric proposed by The School District of Clayton [41]

Further, we adapted the novelty and variety metrics as demonstrated by Linsey et al. [14] for evaluating ideation outcome of our studies. The metrics are defined as follows.

- Variety: The raters were asked to create an exhaustive list of cluster of ideas after thoroughly going through all the mind-maps created by the users. The variety score was then given by the percentage of clusters that was present in the given mind-map.
- Novelty: The novelty score for the mind-maps were calculated by considering the number of other mind-maps

present in the same cluster. That is, lower the number of ideas in cluster, higher the novelty. The following formulae was used for evaluation of the novelty, where N_j is the Novelty score of the j^{th} mind-map, T is the total number of mind-maps, C_i is the number of mind-maps in the i^{th} cluster and n is the number of clusters occupied by the j^{th} mind-map.

$$N_j = \frac{1}{n}\sum_{i=1}^{n}\frac{T - C_i}{T}$$

The inter-raters were given the all the mind-maps and the specific grading rubric. The two raters independently evaluated all the mind-maps for each of these metrics. Further, they were encouraged to discuss and come to a consensus on their grading rubric. The modified values of the metrics were then checked for reliability between the two raters. The Cohen's Kappa value for the metrics Structure, Exploratory, Communication, and Extent of Coverage were found to be in the range of 0.3 - 0.4, showing fair level of agreement and reliability between the two raters [42]. Also the Pearson's correlation between raters for variety and novelty scores was found to be above 0.8 (0.83 and 0.92). This value of correlation coefficient is acceptable based on [42].

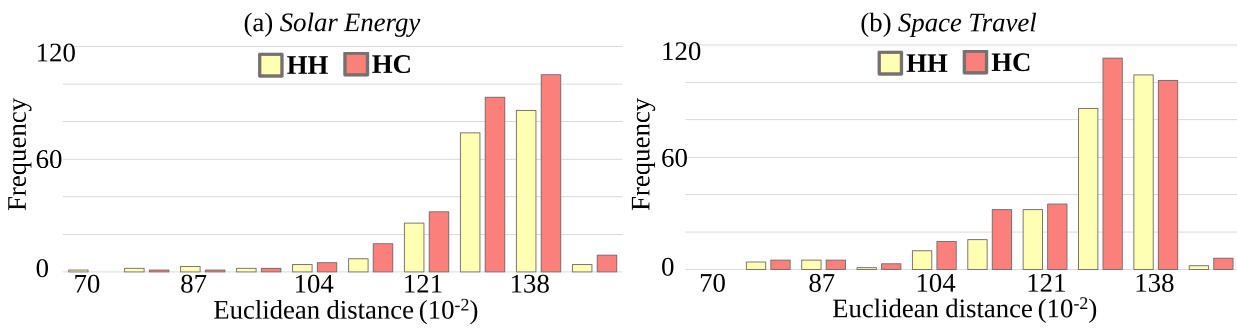

Fig. 7 Distribution of Euclidean distances for word vectors between ideas and corresponding central topic in HH and HC

VI. Results and Discussion

A. Semantic Analysis

For each given topic, we used t-SNE to visualize word embeddings (Figure 8) collected from mind-maps created in both groups of users (**HH** and **HC**) using the pre-trained model from Numberbatch (78 languages, 300 dimensions) [2]. For a given given topic, we observe that most of the generated ideas seem to cluster in proximities of identifiable regions and close to central topics regardless of **HH** or **HC** approach. This indicates that overall, Mini-Map can be comparable to human-human collaborative mind-mapping in terms of the content generation. However, the structure of the mind-maps were quite different in Mini-Map as compared to **HH**; users in **HH** tended to generate more ideas linked with central topics than explore further, and users in **HC** were confined to the standard radial layout. In a way, such non-hierarchical structure of the **HC**s can be used as effective design thinking tools [23].

Using Numberbatch vector space embeddings (Figure 8), we performed semantic similarity comparison by measuring Euclidean distance between word vectors of ideas with its corresponding central topics. There are two main observations to be made regarding the corresponding semantic distance distribution with *Solar Energy* and *Space Travel* (Figures 7). First, the maximum distance in *Solar Energy* recorded for **HH** is 1.49, is the same as recorded for **HC**. In *Space Travel*, the maximum distance is 1.47 for **HH**, whereas **HC** gets a higher score of 1.5. Thus, wherein Mini-Map is able to generate problem-specific content, it is potentially helpful to users while working on open-ended topics. Second, apart from the measured maximum distance which can be attributed to the broadness of one created mind-map, we observe interesting trends in the distribution of distances between word vectors across **HH** and **HC** groups (Figures 7). Ideas from both groups display high frequency in the distance range between 1 to 1.4, however, for **HC**s, the frequency stay higher than **HH**s until the maximum distance ranges observed for both *Solar Energy* and *Space Travel*. This uncovers the potential of Mini-Map which is capable of introducing diversified but still related content to users in creation of mind-maps.

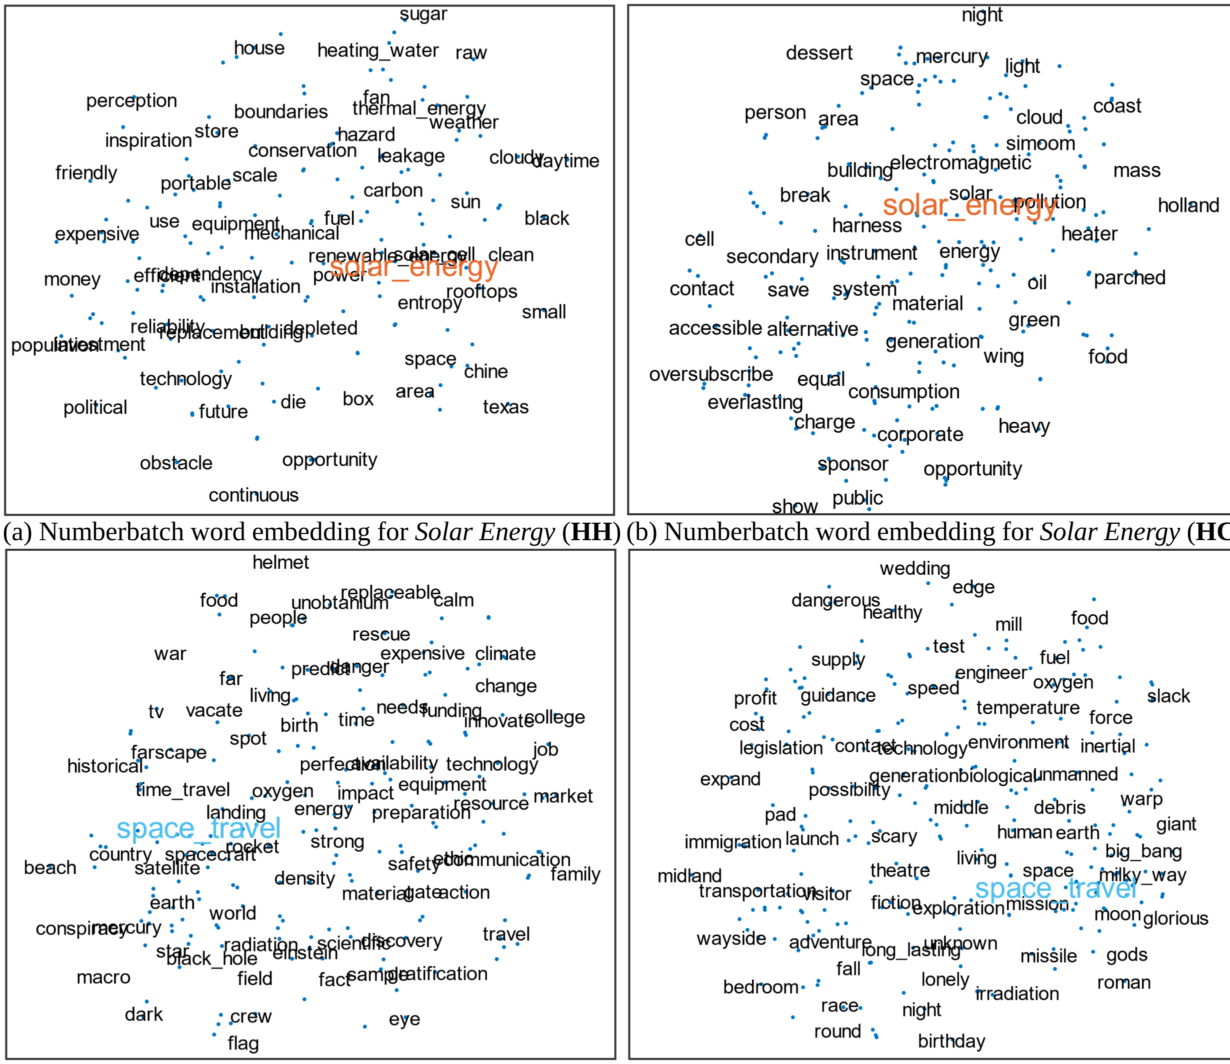

(a) Numberbatch word embedding for *Solar Energy* (**HH**) (b) Numberbatch word embedding for *Solar Energy* (**HC**)

(c) Numberbatch word embedding for *Space Travel* (**HH**) (d) Numberbatch word embedding for *Space Travel* (**HC**)

Fig. 8 Visual representation (t-SNE) of word embedding for *Solar Energy* and *Space Travel* in HH and HC

B. Inter-rater Results

Two-way ANOVA was carried out with two independent variables — type of collaborator and choice of topic. We found that the p-values across the **HH** and **HC** study were almost equal to zero for nearly all metrics (except for novelty whose p-value across **HH** and **HC** was 0.12) indicating a statistically significant difference. However, the p-values across topics were not less than 0.05 showing that the differences were not significant across topic. Results show that the mean of the scores given by the inter-raters for all metrics except novelty was greater in **HC** study compared to **HH** for both the topics (Figure 9). This strongly suggests that Mini-Map helped the users to develop a better mind-maps overall. Specifically, the high values of Structure indicate that the mind-maps created with the Mini-Map interface helped the user create a well-organized map. Perhaps, the target search algorithm might have guided the user to develop a well-structured mind-map by finding the right place to add the node. Another important metric that showed significant improvement is Communication and Extent of Coverage. This highlights how our algorithm for content generation has utilized ConceptNet suitably and generated pertinent and rich vocabulary. The considerable increase in the value of quantity(raw) shows that the Mini-Map algorithm could potentially help the users to generate more ideas within a stipulated time compared to a human collaborator. Also, higher values of variety scores in **HC** suggests that the median-weight algorithm has assisted the user to explore diverse directions. Such differences was not observed in the values of novelty scores. This could mean that although the Mini-Map algorithm cover more number ideas, it was not necessarily unique.

It is to be noted that with small sample of user data, scores for few metrics were not completely normality distributed. However, we carried out the ANOVA for the inter-rater data, since ANOVA is not sensitive moderate deviations from normality. We believe that with greater sample size, we will have a more normally distributed data. However, major insights of our work was exposed in the qualitative part of the study.

Condition	Structure (1-4)	Exploratory (1-4)	Communication (1-4)	Extent of Coverage (1-4)	Quantity (raw)	Variety	Novelty (0-1)
HH Solar Energy	2.43	2.36	2.39	2.21	21.50	66%	0.16
HH Space Travel	2.39	2.18	2.50	2.43	28.29	68%	0.20
HC Solar Energy	3	3.18	2.96	3.04	34.71	78%	0.21
HC Space Travel	3.54	3.36	3.14	3.29	44.43	82%	0.19
Average **HH**	2.41	2.27	2.45	2.32	21.71	67%	0.18
Average **HC**	3.27	3.27	3.05	3.16	35.07	80%	0.20

Fig. 9 The values of various metrics were averaged across topics and type of the collaborator. This table summarizes the mean values of various metrics calculated by the expert raters.

C. User Feedback

1. Qualitative Feedback

We elicited the participant's experience in collaborative mind-mapping using a 7-point likert scale survey questions (Figure 10). In terms of ease of expansion of ideas and thoughts, the results show that Mini-Map was equally conducive compared to a human collaborator. Interestingly, in the **HH** study for the topic *Solar Energy*, around 40% of the participants were not satisfied with the mind-maps they created compared to **HC** study (about 10%). This might potentially be because of the mismatch of ideas that the two users wanted to externalize causing a dissatisfaction in the quality of the mind-map created, from the perspective of the user. Moreover, the response time was totally dependent on the human collaborator in **HH** studies. Given a limited time, delayed responses from the users might have curbed the user to develop a mind-map to their fullest potential — which likely causes a dissatisfaction between the users in the **HH** study.

Majority of the users felt that the collaborator was creative in both **HH** and **HC** study. This shows that the median-based algorithm used for concept generation gave interesting responses commensurate with human-level creativity. Though a bit lower in the **HC** study, majority of the users from both the studies felt that the collaborative mind-mapping would be a better environment for exploring more ideas than individual mind-mapping. These results suggests that Mini-Map is at par with a human collaborator not only in terms of assisting the users to developing greater number of ideas, but also by giving intriguing responses. As per one participant from **HH** study, *"Integration of ideas from the collaborator increased the overall quality of the mind-map"*. Another user from **HC** study stated: "*My collaborator was smart and creative, helped explore ideas - and gave perspectives I hadn't thought of*". Thus, a controlled collaborative setup, like our system, can be particularly helpful in early design stages.

Up to 50% of the users *somewhat disagreed* that the collaborator was not good at organizing their thoughts in the **HC** study. Potential reason for this observation include the difference in approach of making mind-maps of the user compared to the Mini-Map algorithm. This is contrasting to the outcome from the inter-raters claiming that mind-maps created by the **HC** fared significantly well in terms of metrics like Structure and Extent of Coverage. Although our algorithm does not go well with the user approach wise, it guides the user to create a better mind-map overall. This is also corroborated by the fact that majority of the participants were satisfied with the mind-maps they created at the end of the **HC** study.

In **HC** study, one general observations was that the context of the node generated slowly started to depend on the nodes added by the human as time progresses. A participant of **HC** study stated: *"At first it was making too good of thoughts. Then it made some dull contributions"*. So, if the nodes added by the human is well related to the central topic, then the nodes added by the computer maintained the context and gave valuable responses. However, in some cases the connection between the nodes added by human were not explicit. This might not always be identified with the limited semantic data base available, resulting in content generated using median weight method. In **HH** study, there all more than 80% of the users felt that their partner was knowledgeable. However, there were mixed responses from the user about the collaborator's knowledge in the **HC** study. In the topic *Space Travel*, about 75% of the responses were neutral or above. One user stated: *"It was helpful for the topic Space Travel as I had minimal knowledge on that"*. The responses regarding the collaborator's knowledge were slightly more negative for the topic *Solar Energy*. Since this is a more specific topic, there could have been a greater room for diversifying, thereby increasing the possibility of losing context to the idea generated by the median weight method.

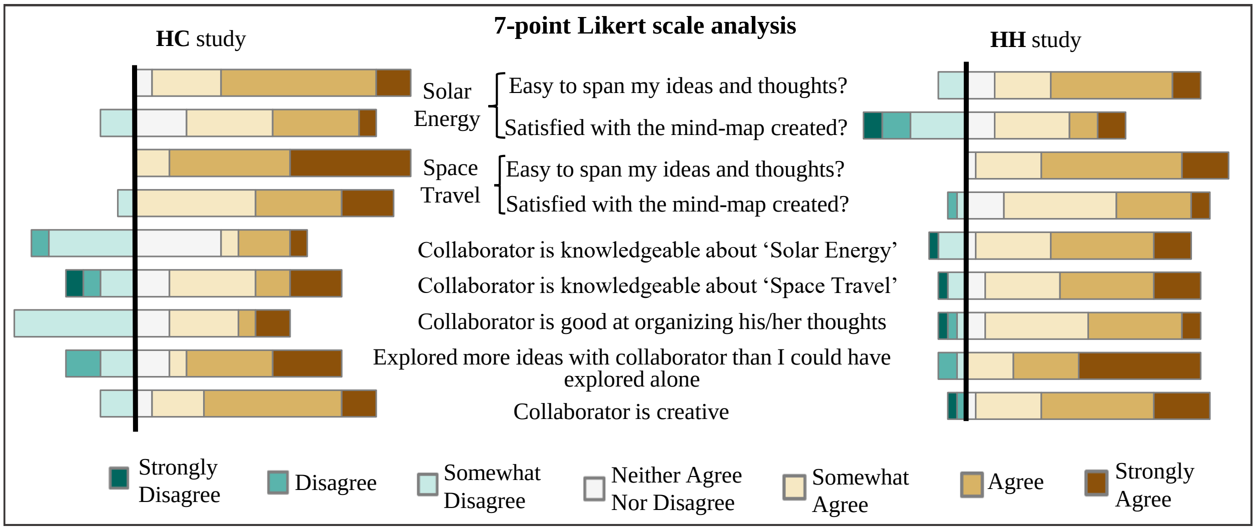

Fig. 10 7-point likert scale feedback from the user study conducted. The brown bars towards the right of the central line indicate positive responses and blue bars to the left indicate negative responses

2. User's Behavioral Perception of the Collaborator

Number of users		Predicted Collaborator	
		Human	Computer
Actual Collaborator	Human	16	12
	Computer	5	9

Fig. 11 This table shows the number of users who predicted their collaborator to be a computer or a human in two types of study conducted.

5 out of 14 users predicted their collaborator to be a human in the **HC** study conducted. Thus, our algorithm seems to have simulated a human-like behavior 35.7% of the time. Interestingly, 12 out of 28 users from the **HH** study predicted their collaborator to be computer. Interesting reasons for predicting Mini-Map to be human includes feedback like *"Smart and creative"* and even *"sly humour"*.

In both **HH** and **HC** study, 35.7% and 42.8% of the participants predicted their collaborator wrongly. Thus, there may exist uncertainty in claiming human-like behavior in Mini-Map algorithm just from this answer. However, open-ended feedback from the users helped us elicit the reason behind their prediction of the type of collaborator. It also helped us to develop understanding of how Mini-Map can simulate the useful aspects of human-like behaviors and ways in which Mini-Map can prove to be more helpful than a human collaborator.

D. Median-weighted Node vs Relation Node

Adding relationship node was mainly found to be helpful for categorizing the central ideas into branches. For example, user responses suggests that nodes like *'Space Travel has...', 'Requirement'*, in the topic *Space Travel* and nodes like *'Capable of...'* in *Solar Energy* specifically gave fundamental directions to think about the given central idea. Additionally, these nodes could have assisted the users to get started with mind-mapping analogous to initiating a discussion to increase a person's engagement. Moreover, one of the user's statement: *"I felt some part like a conversation"* might allude to the relationship type nodes.

E. Limitations

While having capabilities to explore and generate reasonable ideas around a central topic, the current status of Mini-Map also has several limitations. First of all, at times, Mini-Map generate contents that are repetitive to context in current map, synonyms, or overly sophisticated words that could not be interpreted by a common individual. Also, there were instances where the content generation algorithm did not generate nodes directly related with the central topic. Secondly, some median-weighted nodes added were not particularly helpful for further exploration. Rather, it was analogous to a question and answer. For example, the node *'long been a dream of mankind'* was added to the relation node *'Space Travel has...'*. Thirdly, our algorithm is dependent on the limited data available in ConceptNet which is largely a general relation dependent knowledge base. Thus, it has limited domain-specific knowledge.

VII. Conclusions and Future Directions

At its core, Mini-Map presents a digital mind-mapping work-flow for co-generating ideas with an intelligent agent as a collaborating partner. Technically, we make two main contributions in the work-flow of the Mini-Map. First, we incorporated guidelines on how to mind-map for identifying target nodes. Second, we demonstrated a relation-dependent and path-dependent method to extract relevant content from ConceptNet to enable mind-map evolution. The user study conducted consists of two groups of participants (**HH** and **HC**) to understand and evaluate the processes and outcomes in collaborative mind-mapping. We found that Mini-Map could perform to the level of a human collaborator, in terms of assessment from both semantic (Numberbatch embedding) and subjective (structure, exploratory, communication, extent of coverage, variety, novelty) perspectives. In addition, even though our primary purpose was not to show a system that fooled the user in believing that the collaborator was human, our system was able to . This could be a promising prospect for future mixed initiative systems for novice users to learn and develop skills in unstructured design tasks.

There are several promising research directions that we envisage continuing with research. Our goal in the future is to improve target search algorithm in Mini-Map by incorporating brain computer interfaces (BCI) or a cognitive model. Such approach can provide us with information on identification of user attention level and preferences for a certain contents in the mind-map, hence adjust Mini-Map's behavior accordingly. In addition to implementing a reasoning model, we are also interested in making our study crowd-sourcing since numerous time was taken in collecting user data in a controlled experimental setup. Also, crowd-sourcing helps in collecting significant amount of data for a comprehensive analysis. Moreover, we look forward to develop a database specific to design ideation that would bring the most out of the Mini-Map workflow.

The current state-of-art in human automation interactions (Siri, Cortana, Alexa, and Google Assistant) are modeled based on the metaphor of an intelligent assistant — these are solution-oriented systems that provide answers to reasonably well-formulated problems. While there have been a few works on mixed-initiative design for specific domains such as AI based game development [43], the context of the game is fixed and the designer is only being helped with the detailing of the game rather than finding out what the game should do in the first place. We believe that the next major advances in mixed-initiative interactions should focus on systems that help find good problems. We believe that our work takes a major step toward that goal by demonstrating human-computer collaboration for highly unstructured tasks such as mind-mapping.

References

[1] Speer, R., and Havasi, C., "ConceptNet 5: A large semantic network for relational knowledge," *The People's Web Meets NLP*, Springer, 2013, pp. 161–176.

[2] Speer, R., Chin, J., and Havasi, C., "ConceptNet 5.5: An Open Multilingual Graph of General Knowledge." *AAAI*, 2017, pp. 4444–4451.

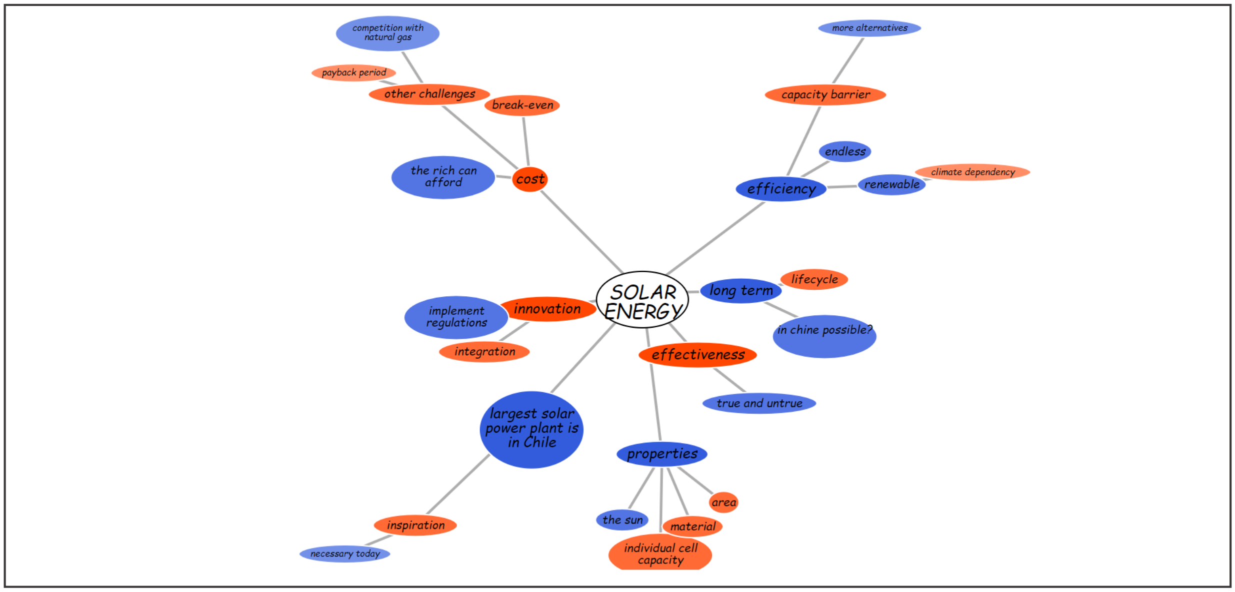

(a) Mind-map created with *Solar Energy* in **HH** group

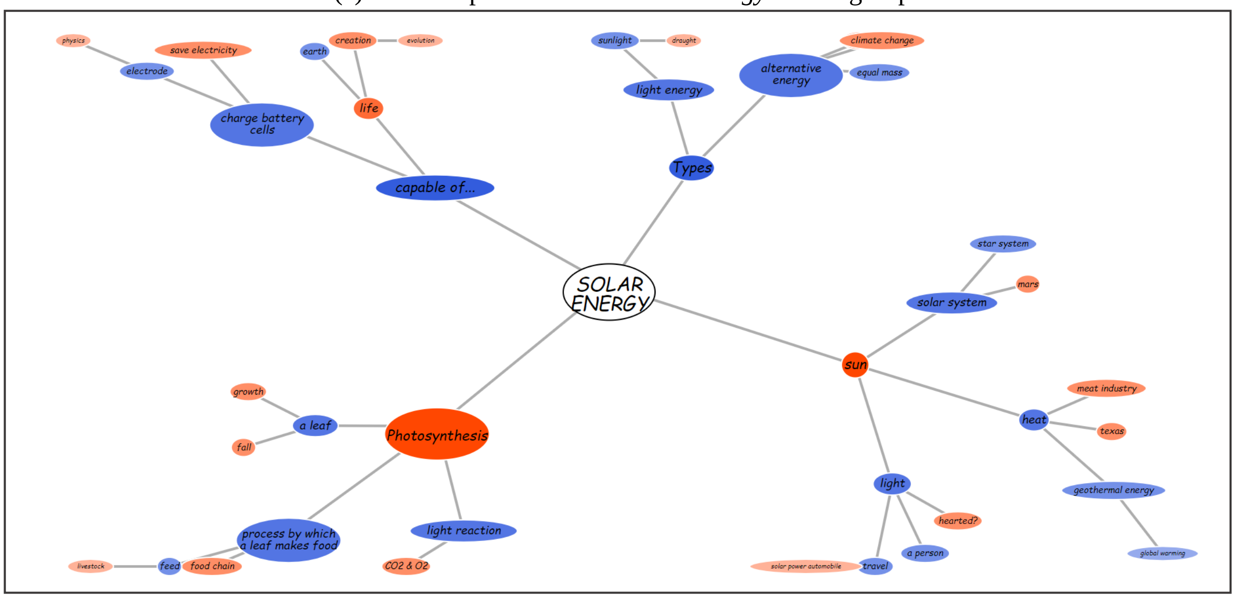

(b) Mind-map created with *Solar Energy* in **HC** group

Fig. 12 Sample user created mind-maps with central topic *Solar Energy* in (a) HH and (b) HC. In (b), orange nodes represent user added nodes, and blue nodes represent Mini-Map added nodes.

[3] Liu, H., and Singh, P., "ConceptNet—a practical commonsense reasoning tool-kit," *BT technology journal*, Vol. 22, No. 4, 2004, pp. 211–226.

[4] Havasi, C., Speer, R., and Alonso, J., "ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge," *Recent advances in natural language processing*, Citeseer, 2007, pp. 27–29.

[5] Lin, H., and Faste, H., "Digital Mind Mapping: Innovations for Real-time Collaborative Thinking," *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA, 2011, pp. 2137–2142. doi:10.1145/1979742.1979910, URL `http://doi.acm.org/10.1145/1979742.1979910`.

[6] "FreeMind - free mind mapping software," `http://freemind.sourceforge.net/wiki/index.php/Main_Page`, 2018.

[7] "mindjet," `https://www.mindjet.com/`, 2018.

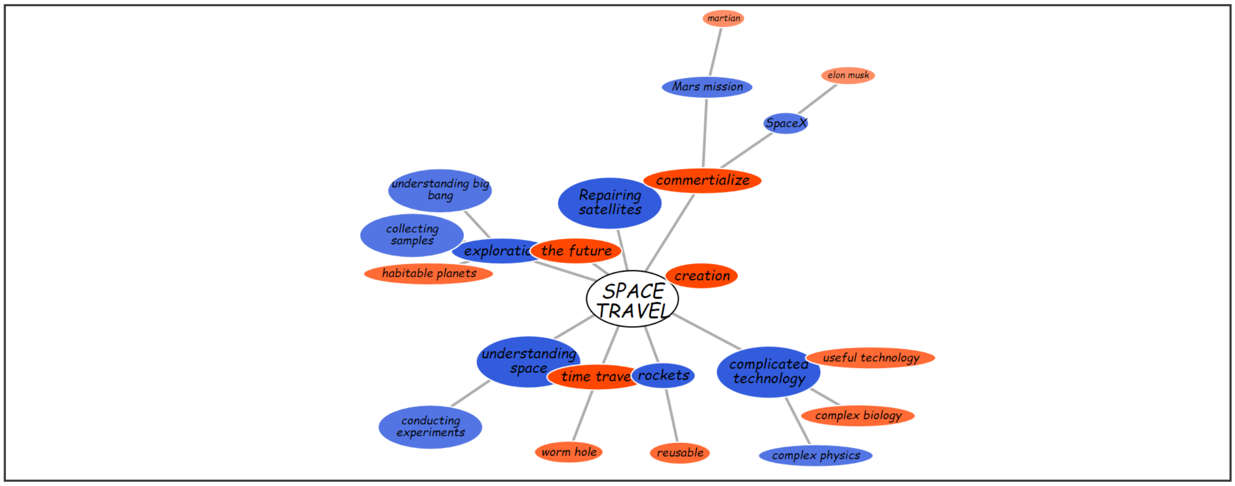

(a) Mind-map created with *Space Travel* in **HH** group

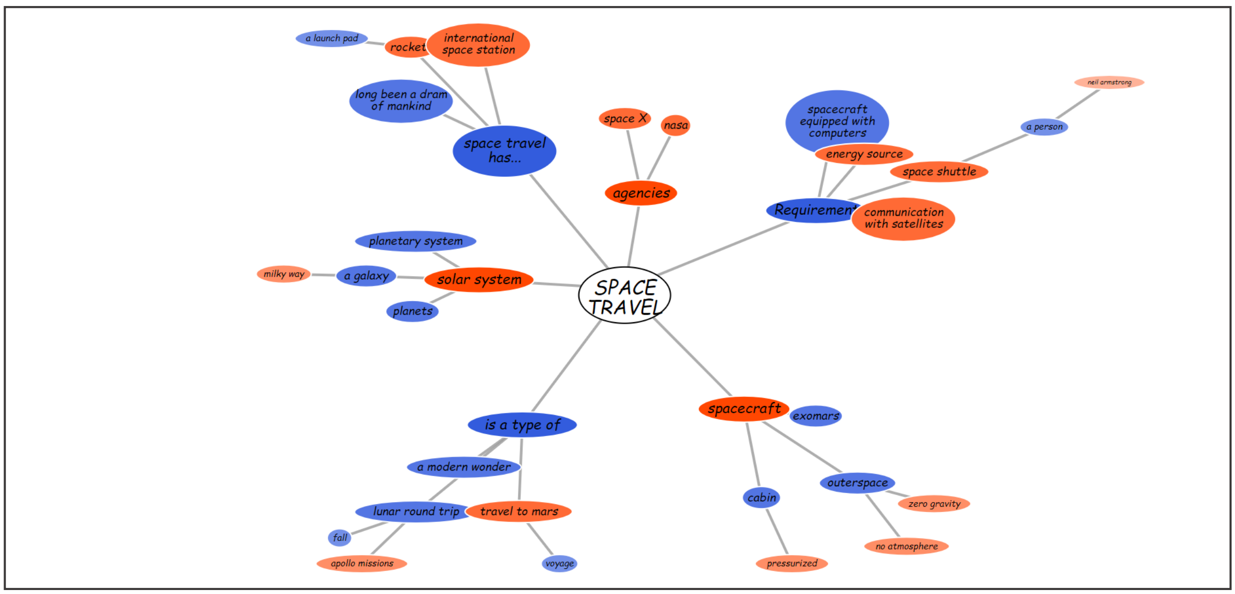

(b) Mind-map created with *Space Travel* in **HC** group

Fig. 13 Sample user created mind-maps with central topic *Space Travel* in (a) HH and (b) HC. In (b), orange nodes represent user added nodes, and blue nodes represent Mini-Map added nodes.

[8] "Mind Mapping Software: Brainstorm Online," `https://www.mindmeister.com/`, 2018.

[9] "Online Diagram and Flowchart software," `https://cacoo.com/`, 2018.

[10] "Diagram Maker," `https://creately.com/`, 2018.

[11] "Online Diagram Software and Visual Solution," `https://www.lucidchart.com/`, 2018.

[12] "Tinkmap Visual Thesaurus," `https://www.visualthesaurus.com/`, 2018.

[13] Yannakakis, G. N., Liapis, A., and Alexopoulos, C., "Mixed-initiative co-creativity." *FDG*, 2014.

[14] Linsey, J. S., Clauss, E., Kurtoglu, T., Murphy, J., Wood, K., and Markman, A., "An experimental study of group idea generation techniques: understanding the roles of idea representation and viewing methods," *Journal of Mechanical Design*, Vol. 133, No. 3, 2011, p. 031008.

[15] Willis, C. L., and Miertschin, S. L., "Mind Tools for Enhancing Thinking and Learning Skills," *Proceedings of the 6th Conference on Information Technology Education*, ACM, New York, NY, USA, 2005, pp. 249–254. doi:10.1145/1095714.1095772, URL http://doi.acm.org/10.1145/1095714.1095772.

[16] Willis, C. L., and Miertschin, S. L., "Mind Maps As Active Learning Tools," *J. Comput. Sci. Coll.*, Vol. 21, No. 4, 2006, pp. 266–272. URL http://dl.acm.org/citation.cfm?id=1127389.1127438.

[17] Kommers, P., and Lanzing, J., "Students' Concept Mapping for Hypermedia Design: Navigation Through World Wide Web (WWW) Space and Self-assessment," *J. Interact. Learn. Res.*, Vol. 8, No. 3-4, 1997, pp. 421–455. URL http://dl.acm.org/citation.cfm?id=313722.313732.

[18] Faste, R., "mindmapping," http://www.fastefoundation.org/publications/mind_mapping.pdf, ????

[19] Holland, B., Holland, L., and Davies, J., "An investigation into the concept of mind mapping and the use of mind mapping software to support and improve student academic performance." University of Wolverhampton, 2004.

[20] Burke, L. A., and Miller, M. K., "Taking the mystery out of intuitive decision making," *The Academy of Management Executive*, Vol. 13, No. 4, 1999, pp. 91–99.

[21] Zampetakis, L. A., Tsironis, L., and Moustakis, V., "Creativity development in engineering education: The case of mind mapping," *Journal of Management Development*, Vol. 26, No. 4, 2007, pp. 370–380.

[22] Isaksen, S. G., Dorval, K. B., and Treffinger, D. J., *Creative approaches to problem solving: A framework for change*, Kendall Hunt Publishing Company, 2000.

[23] Kokotovich, V., "Problem analysis and thinking tools: an empirical study of non-hierarchical mind mapping," *Design studies*, Vol. 29, No. 1, 2008, pp. 49–69.

[24] Buzan, T., *The ultimate book of mind maps: unlock your creativity, boost your memory, change your life*, HarperCollins UK, 2006.

[25] Buisine, S., Besacier, G., Najm, M., Aoussat, A., and Vernier, F., "Computer-supported Creativity: Evaluation of a Tabletop Mind-map Application," *Proceedings of the 7th International Conference on Engineering Psychology and Cognitive Ergonomics*, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 22–31. URL http://dl.acm.org/citation.cfm?id=1784197.1784201.

[26] Faste, H., and Lin, H., "The Untapped Promise of Digital Mind Maps," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 2012, pp. 1017–1026. doi:10.1145/2207676.2208548, URL http://doi.acm.org/10.1145/2207676.2208548.

[27] Orehovački, T., Granić, A., and Kermek, D., "Exploring the quality in use of Web 2.0 applications: the case of mind mapping services," *International Conference on Web Engineering*, Springer, 2011, pp. 266–277.

[28] Abdeen, M., El-Sahan, R., Ismaeil, A., El-Harouny, S., Shalaby, M., and Yagoub, M. C. E., "Direct automatic generation of mind maps from text with M2Gen," *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*, 2009, pp. 95–99. doi:10.1109/TIC-STH.2009.5444360.

[29] Elhoseiny, M., and Elgammal, A., "English2MindMap: An Automated System for MindMap Generation from English Text," *2012 IEEE International Symposium on Multimedia*, 2012, pp. 326–331. doi:10.1109/ISM.2012.103.

[30] Elhoseiny, M., and Elgammal, A., "Text to multi-level MindMaps," *Multimedia Tools and Applications*, Vol. 75, No. 8, 2016, pp. 4217–4244. doi:10.1007/s11042-015-2467-y, URL https://doi.org/10.1007/s11042-015-2467-y.

[31] Kudelić, R., Konecki, M., and Maleković, M., "Mind map generator software model with text mining algorithm," *Proceedings of the ITI 2011, 33rd International Conference on Information Technology Interfaces*, 2011, pp. 487–494.

[32] Deterding, S., Hook, J., Fiebrink, R., Gillies, M., Gow, J., Akten, M., Smith, G., Liapis, A., and Compton, K., "Mixed-Initiative Creative Interfaces," *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA, 2017, pp. 628–635. doi:10.1145/3027063.3027072, URL http://doi.acm.org/10.1145/3027063.3027072.

[33] Liapis, A., Yannakakis, G. N., and Togelius, J., "Sentient Sketchbook: Computer-aided game level authoring." *FDG*, 2013, pp. 213–220.

[34] Bang, H., Virós Martin, A., Prat, A., and Selva, D., "Daphne: An Intelligent Assistant for Architecting Earth Observing Satellite Systems," *2018 AIAA Information Systems-AIAA Infotech@ Aerospace*, 2018, p. 1366.

[35] Kotov, A., and Zhai, C., "Tapping into Knowledge Base for Concept Feedback: Leveraging Conceptnet to Improve Search Results for Difficult Queries," *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*, ACM, New York, NY, USA, 2012, pp. 403–412. doi:10.1145/2124295.2124344, URL `http://doi.acm.org/10.1145/2124295.2124344`.

[36] Miller, G. A., "WordNet: a lexical database for English," *Communications of the ACM*, Vol. 38, No. 11, 1995, pp. 39–41.

[37] Gilon, K., Chan, J., Ng, F. Y., Liifshitz-Assaf, H., Kittur, A., and Shahaf, D., "Analogy Mining for Specific Design Needs," *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, 2018, pp. 121:1–121:11. doi:10.1145/3173574.3173695, URL `http://doi.acm.org/10.1145/3173574.3173695`.

[38] Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J., "Distributed Representations of Words and Phrases and Their Compositionality," *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, Curran Associates Inc., USA, 2013, pp. 3111–3119. URL `http://dl.acm.org/citation.cfm?id=2999792.2999959`.

[39] Pennington, J., Socher, R., and Manning, C. D., "GloVe: Global Vectors for Word Representation," *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543. URL `http://www.aclweb.org/anthology/D14-1162`.

[40] Maaten, L. v. d., and Hinton, G., "Visualizing data using t-SNE," *Journal of machine learning research*, Vol. 9, No. Nov, 2008, pp. 2579–2605.

[41] "Scoring Rubric for Mind Maps," `https://www.claytonschools.net/cms/lib/MO01000419/Centricity/Domain/206/Mind%20Map%20Rubric.pdf`, 2008.

[42] Clark-Carter, D., *Doing Quantitative Psychological Research: From Design to Report*, Taylor & Francis, Inc., UK., 1997.

[43] Baldwin, A., Dahlskog, S., Font, J. M., and Holmberg, J., "Mixed-initiative procedural generation of dungeons using game design patterns," *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, 2017, pp. 25–32. doi:10.1109/CIG.2017.8080411.