Entrepreneurship: A New Frontier in a Computational Science of Creativity

Ashok Goel, Mukundan Kuthalam, Sung Jae Hong, Keith McGregor
Design Intelligence Laboratory, School of Interactive Computing,
Georgia Institute of Technology, Atlanta, USA

Abstract

We present entrepreneurship as a new frontier for developing a computational science of creativity. Entrepreneurship often requires a large upfront investment, but customer feedback typically is delayed, and while the costs of failure can be considerable, the gains from success too can be significant. Given its importance, how can we help novice entrepreneurs learn about entrepreneurship? We analyze the role human coaches play in mentoring novice entrepreneurs by asking critical questions to help generate business models. We propose that virtual coaches may augment the learning of novice entrepreneurs. We describe a preliminary experiment in designing a virtual coach named Errol for learning about entrepreneurship. When a startup team creates an initial business model, Errol uses semantic and lexical analyses to ask questions about the model, leading to model revision. Our experiment indicates that creativity may emerge out of the interactions between the virtual coach and a startup team, and that this human-computer co-creativity may accelerate the process by which a novice entrepreneurs can learn to create intermediate-level business models.

Keywords: Entrepreneurship, Human-Computer Co-Creativity, Question Asking, Virtual Coaches.

Introduction

The field of computational creativity seeks to develop a computational science of creativity, including computational theories of creativity, techniques for realizing artificial creativity, and tools for supporting human creativity. It addresses creativity in a variety of domains such as the arts, design, literature, science, etc. Our first goal in this position paper is to introduce entrepreneurship as a new frontier for research on computational creativity. Entrepreneurship often requires a large upfront investment, the evaluation through customer feedback typically is delayed, the chances for success are low, and the costs of failure can be high. Yet, when an enterprise succeeds, it can make a difference in the world and the gains can be significant.

Given the importance of entrepreneurship to the economy, the workforce, and the society, many academic institutions around the world are expanding educational programs in creating startups. Thus, our second goal in this article is to present a high-level information-processing account of teaching and learning about entrepreneurship in which experienced serial entrepreneurs act as coaches to novice entrepreneurs. For example, the US NSF's Innovation Corps (I-Corps) (Huang-Saad, et al. 2017) runs educational programs for scientists and engineers to learn how to instigate such startups. In these programs, human coaches teach fundamentals of entrepreneurship such as customer discovery and business model generation. Much of the teaching uses the Socratic method in which the coaches ask critical questions to encourage reflection by the novice entrepreneurs. Based in part on the questions asked by coaches, novice entrepreneurs learn to build better business models and often make a pivot in their value propositions and customer segments. The creativity of the business models emerges out of these interactions between the coaches and the novice entrepreneurs.

Third, we posit that virtual coaches may amplify the reach of human coaches and augment the learning of novice entrepreneurs. Serial entrepreneurs who are willing to act as coaches to novice entrepreneurs are not easy to find. Further, human coaches often have biases. Given that the questions human coaches ask in early phases of business model generation tend to have well-established patterns, a virtual coach may ask similar questions based on commonly accepted norms. This may help reduce bias while accelerating the process of learning to generate intermediate-level business models. Virtual coaches have the added benefits that they can be used by anyone, anytime, anyplace. Thus, they can support learning how to create a startup on demand and at scale, and thereby help foster a culture of entrepreneurship.

Finally, we briefly describe a preliminary experiment in developing a "proof-of-concept" virtual coach named Errol for teaching entrepreneurship. When a startup team creates an initial business model, Errol uses semantic and lexical analyses of the entries in the model to ask questions. This leads to reflection by the startup team, and results in iterative revision and refinement of the model. By attempting to categorize and correct the errors that novices typically make, Errol seeks to accelerate the process by which a novice startup team can start creating intermediate-level business models. Errol is an experiment in realizing human-computer co-creativity in entrepreneurship with the creativity of the business models emerging out of its interactions with novice entrepreneurs.

The Domain of Entrepreneurship

There is broad consensus in research on computational creativity on several aspects of creativity. For example, creativity requires both a producer and a receiver; creative products are novel, useful and non-obvious; and creativity lies on a spectrum (e.g., Besold, Schorlemmer, and Smaill 2015; Boden 1990; Sternberg 1999; Veale and Cardoso 2020). However, research on computational creativity continues to explore new domains of creativity that are often characterized by new dimensions of the creativity spectrum. For example, in early research on AI in design, Goel and Chandrasekaran (1991) viewed design as a case-based process. They described a spectrum of creativity that starts at one end with the solution to a new design problem being identical to that for a previous problem; moves to the solution to the new design problem differing from that of a previous problem only in the values of parameters of a component; then to the solution to the new design problem differing from that of a previous problem through the substitution of one component by another; next to the solution to the new problem differing from that of a previous problem through the addition or deletion of a component, and so on. In contrast, more recently Fitzgerald, Goel, and Thomaz (2017) have viewed robot creativity as a process of embodied analogical reasoning. They describe a different spectrum of creativity that starts on one end with the new task being identical to a task familiar to the robot; moves to the new task differing from the familiar task in the parameters of an object; then to the new task differing from the familiar task in the object themselves; next to the new task differing from the familiar task in the relationships among the objects, and so on.

Characterization of creativity in literature and the arts typically requires consideration of additional dimensions (Candy and Edmonds 2002; Cohen 1995; Cope 2005; Pearce and Wiggins 2014; Perez y Perez and Sharples 2001; Veale 2012). Such additional dimensions may pertain to the problem definition because the problem may be ill-defined, and/or the evaluation criteria for the criteria for evaluation may be unknown in advance or difficult to operationalize. The point here again is that each new domain of creativity appears to introduce new dimensions of the spectrum of creativity

We posit entrepreneurship as a new frontier in creativity, one that offers new challenges both for building a computational theory of creativity and for developing interactive tools for supporting human creativity. As with creativity in general, creativity in entrepreneurship requires both a producer and a receiver, and the products of creativity are at least initially novel, useful, and non-obvious to the intended receiver. Further as with creativity in literature and the arts, problems in entrepreneurship are ill-defined and the evaluation criteria are unknown in advance or difficult to operationalize.

In addition, creativity in entrepreneurship is defined by a large upfront financial investment and much delayed evaluation through customer feedback. Further, unlike many other creative domains, the chances of success in entrepreneurship typically are low and the economic costs of failure can be high. However, when an enterprise succeeds, it can make a real difference to the lives of the receivers and result in the significant financial gains to the producers as well as economic benefits to the community of producers and receivers as whole. These additional dimensions make entrepreneurship a unique and challenging domain because the margin for errors is very small - much smaller than in many other domains of creativity.

Learning about Entrepreneurship

Entrepreneurship is critical not only to employment, but also to the economy and the society of a country (Carland & Carland 2004). Thus, many institutions around the world are expanding and strengthening their educational programs in entrepreneurship (Kuratko 2004) such as the CREATE-X program at Georgia Tech (Forest et al. 2021). We will use the I-Corps program for academic scientists and engineers as a case study to develop a high-level information-processing account of the teaching and the learning in entrepreneurship education.

The I-Corps program focuses on startup creation by bringing the rigor of engineering science to entrepreneurship. Thus, the novice entrepreneurs are encouraged to think of entrepreneurship in terms of the scientific process: hypothesis generation, experiment design, data collection, hypothesis revision and refinement, and so on. A startup's business model articulates and elaborates on its business hypotheses. In the I-Corps program, serial entrepreneurs act as human coaches to novice entrepreneurs, teaching fundamentals of entrepreneurship such as customer discovery and business model generation. Much of the teaching uses the Socratic method (Paul & Elder 2007) in which the coaches rarely give direct advice; instead, they ask critical questions to encourage reflection by the novice entrepreneurs. Based in part on the questions asked by coaches, novice entrepreneurs learn to build better business models and often revise their hypotheses by making a pivot in their value propositions and customer segments. The creativity of the business models

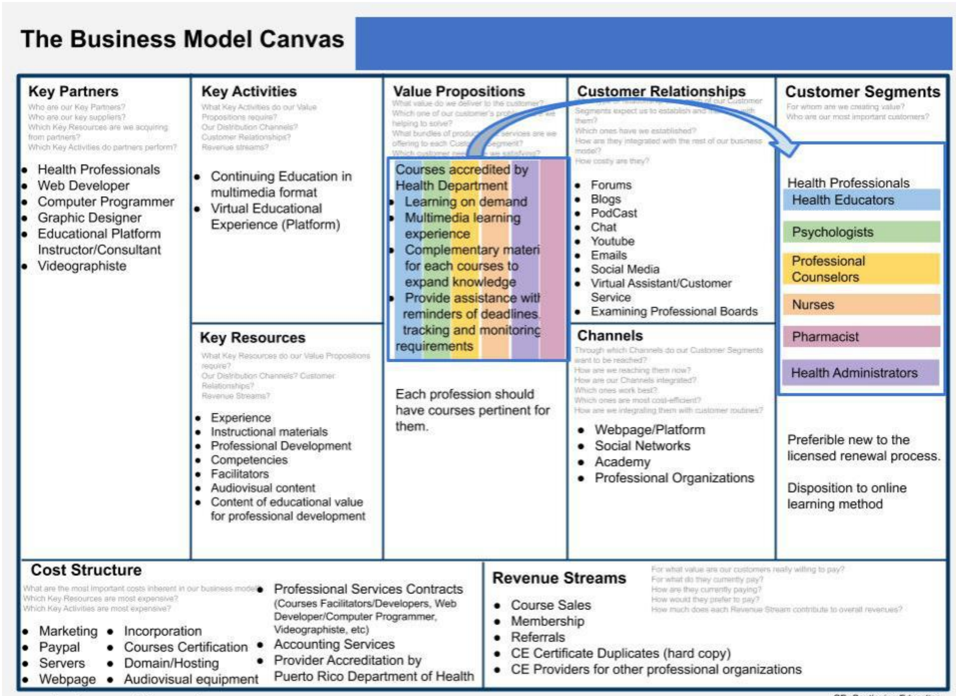

Figure 1: The business model canvas of a team that is creating a product for the continued education of healthcare professionals. The BMC is "complete" in the sense that all the fields are filled in, but the current iteration of Errol would concern itself with just the customer segments and value propositions (the entries with colored backgrounds).

emerges out of the interactions between the coaches and the startup team.

Business Model Canvas

The I-Corps program uses Osterwalder & Pigneur's (2010) Business Model Canvas (BMC) for articulating and elaborating a startup's business hypotheses. A BMC is a modeling template that simplifies the description of a business into three major categories and nine sections within them: Desirability (Customer Segments, Value Propositions, Customer Relationships, Channels), Feasibility (Key Activities, Key Resources, Key Partners), and Viability (Cost Structure, Revenue Streams). Figure 1 provides a snapshot of an evolving business model.

The BMC affords would-be entrepreneurs cognitive support for articulating, elaborating, sharing, and critiquing their hypotheses about the value their potential business adds to some market, the targeted customer segments, and several other components necessary to spawn their startups. Yet developing a good business model using a BMC also requires getting thorough and highly specialized feedback on what is good and what can be improved. Thus, I-Corps coaches provide iterative feedback on a startup's BMC, with the simultaneous intent of enlightening the novice entrepreneur as well as enabling them to provide the right products and services to the right markets, thereby optimizing their potential for success.

Human Coaching

We have extensively consulted with some of the human coaches in the I-Corps program; indeed, the last author on this paper (KM) is a lead instructor in the program. We have also observed human coaching in the I-Corps program in practice; the first author (AG) has twice taken the program to learn about spawning startups. In addition, we have analyzed coaches' interactions with novice entrepreneurs in the teaching of the I-Corps Puerto Rico cohort in 2018. Our analysis indicates that the coaches simultaneously are trying to understand the developing business model on a given BMC as a whole and determining what the various segments on the BMC are lacking in content, structure, and relationships. We found that coaches then formulated questions designed to help the novice entrepreneurs become aware of the relationships described on a BMC into classes of errors to help the novice entrepreneurs become aware of the errors and encourage them to reflect on and seek to address the errors.

While the I-Corps program may be considered a success as evidenced by its rapid expansion over the last decade (Huang-Saad et al. 2017), education in entrepreneurship in general faces several obstacles. The first problem is scaling up the teaching and learning in the I-Corps program to other formal and informal programs in entrepreneurship education. However, serial entrepreneurs who can act as human coaches and provide feedback on business models of novice entrepreneurs are not easy to find. Thus, scaling the teaching and learning in I-Corps to other entrepreneurship programs and startup incubators is difficult. Additionally, even within an educational program, on-demand access to human coaches is limited because of constraints on their available time. Thus, from the perspective of novice entrepreneurs, they do not have adequate access to human coaches who would give them critical feedback on their business models, affording faster improvements to their business modeling practices.

Further, the teaching skills required to provide feedback that strictly adheres to the Socratic method of critical questioning without offering consultative advice is not very common among human coaches. Finally, human coaches have not only many of the typical cognitive biases, but also opinions based on their personal experiences with entrepreneurship. These opinions can make their feedback on a business model proposed by a startup team subjective, biased and skewed.

Virtual Coaches

We posit that virtual coaches may amplify the reach of human coaches and augment the learning of the novice entrepreneurs. The early stages of developing a business model follow a repetitive pattern of common errors by most novice entrepreneurs. Thus, human coaches give the same types of feedback to startup after startup, cohort after cohort. Given that the feedback in the early stages tends to have well-established patterns based on commonly accepted norms, a virtual coach may ask similar questions. This may help reduce bias in the feedback while accelerating the process of learning to generate intermediate-level business models. Once a novice startup team has progressed to an intermediate level of expertise, human coaches may take over from the virtual coach and help the team in reaching advanced levels of expertise. Virtual coaches have the added benefits that they can be used by anyone, anytime, anyplace. Thus, they can support

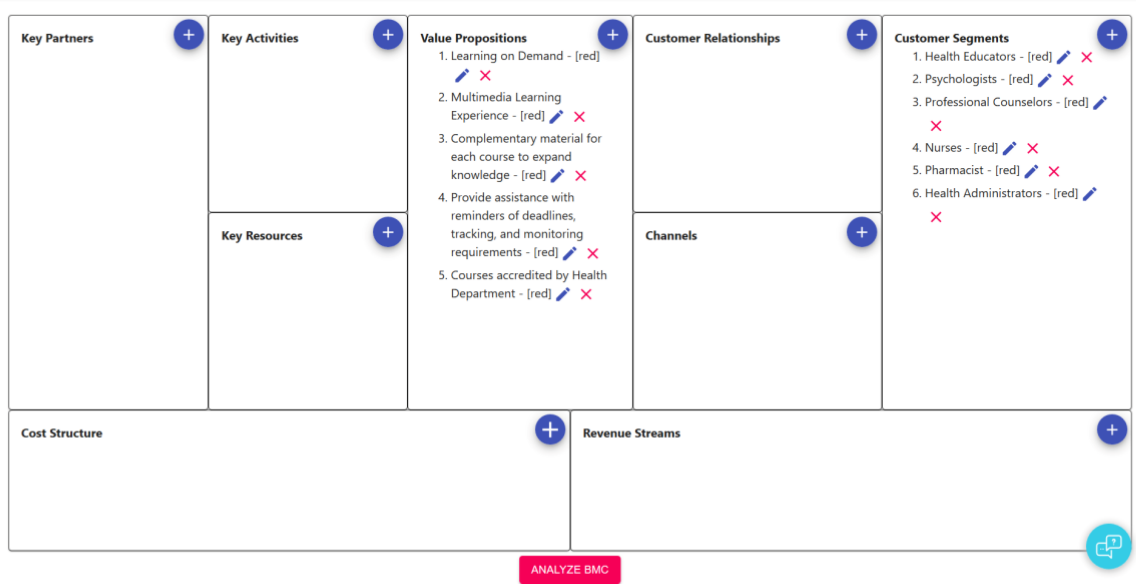

Figure 2: The user interface to Errol. Since Errol currently addresses only value propositions and customer segments, other fields in the BMC have been left blank. The pencil icons allow users to edit entries while the x's are for deletions. The colors, like [red], indicate an underlying relationship between customer segments and value propositions. In this example, all items tagged [red] are related to each other.

learning about entrepreneurship on demand and at scale.

It is important to note that the virtual coach need not, indeed should not, precisely mimic a human coach; instead, the goal is to simulate the experience of a human coach reviewing a business model represented on a BMC and giving directed feedback by asking pointed questions, thereby providing guidance that novice learners of entrepreneurship may find useful in improving their business model. Instead of mimicking a (possibly biased) human coach, a virtual coach should try to address typical novice errors in business model generation.

More broadly, a virtual coach should act like a "Sounding Board" (Schank & Cleary 1995) to the startup team. The key to addressing learning in an open-ended problem such as design or entrepreneurship is to ask the right questions. The startup team likely will know more about the given problem than its coach. Thus, the startup team should maintain control and the virtual coach should only supply the team with questions that it believes are appropriate to pursue. The virtual coach might not know much about the problem at hand; all it will know are what types of questions are useful to ask and how to present the questions in a sensible order. A passive learner may choose to follow the path suggested by the virtual coach, but a more active learner may make a pivot based in part on the questioning, revise the business model, and thereby change the course of questioning. The creativity in the final business model will emerge from the interactions between the virtual coach and the startup team.

Errol: Teaching by Asking Questions

Errol is a preliminary, "proof-of-concept" virtual coach that captures a few of the essential characteristics of an idealized startup virtual coach. Errol focuses on giving feedback by asking questions a business model captured on a BMC, simulating the experience of receiving expert Socratic guidance to improve a business model. Errol takes as input the business model expressed on the BMC and generates questions as output that either ask the learner to clarify the contents of the business model or seek to address the errors common to novice entrepreneurs. By adhering strictly to providing feedback only in the Socratic method, Errol invites the startup team to reflect on its own reasoning and to deliberate on its hypotheses and the assumptions underlying them.

While all nine elements of BMC are relevant to the model of the emerging business, the current version of Errol focuses on the Desirability aspect of the proposed startup, and expressly upon the two most important sections and the relationships between them: Value Propositions and Customer Segments. These two sections are the primary drivers that determine the

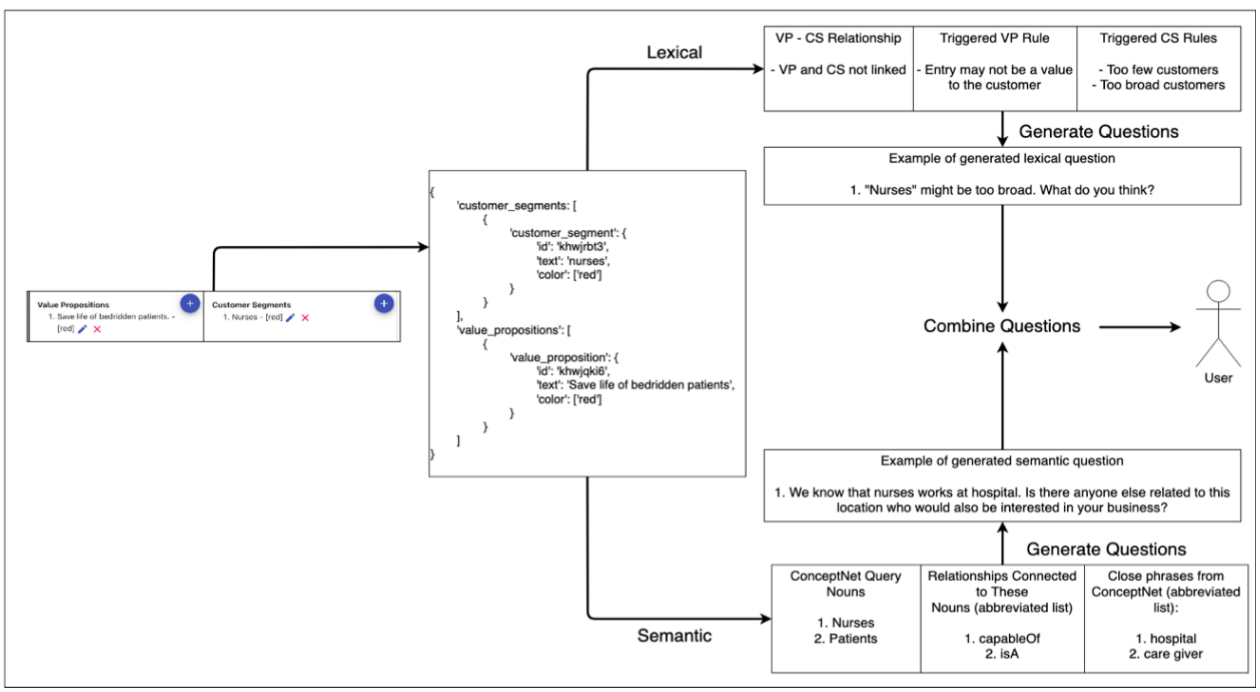

Figure 3: A graphical representation of Errol's processing. Errol extracts from the user's business model value propositions, customer segments, and relationships between them, and provides them to both lexical and semantic processing. The semantic processing focuses on the nouns in each of entries, while the lexical processing goes straight to rules.

Questions Generated by Customer Segment Triggers

Customer Segments	Generated Questions
Health Educators, Psychologists, Professional Counselors, Nurses, Pharmacist, Health Administrators	Overall, could you try to be more specific about who you are trying to target with "Health Educators"?
Health Educators, Psychologists, Professional Counselors, Nurses, Pharmacist, Health Administrators	Can you expand on "Health Educators"?

Questions Generated by Value Proposition Triggers

Value Proposition	Generated Question
Provide assistance with reminders of deadlines, tracking, and monitoring requirements	Would you mind expanding on some of these shortforms in "Provide assistance with reminders of deadlines, tracking, and monitoring requirements"?

Remaining Generated Questions

We know that health is related to being. I would suggest incorporating this idea into your BMC.
We know that counselors are capable of give advice. I would suggest incorporating this idea into your BMC.
We know that nurses are capable of care for patients. I would suggest incorporating this idea into your BMC.
We know that nurses works at a hospital. Is there anyone else related to this location who would also be interested in your business?
We know that nurses are at location hospital. I would suggest incorporating this idea into your BMC.
We know that nurses is a care givers. I would suggest incorporating this idea into your BMC.
We know that pharmacist is a person. I would suggest incorporating this idea into your BMC.
We know that health is related to well. I would suggest incorporating this idea into your BMC.
Is every customer relevant to your business? Have you segmented your customers too much?

Table 1: Questions generated for BMC 709.

remainder of a business model. Customer Segments (CS) declare which customers will be served by the business while Value Propositions (VP) address what value the business will add to satisfy the needs of those Customer Segments. Figure 2 illustrates the user interface to Errol using BMC 709 as an example.

Information Processing in Errol

In this position paper, we describe the information processing in Errol only briefly; Goel et al. (2020) provides more details. Figure 3 illustrates the flow of information in Errol. The entries in the BMC on the left of the figure are read, tagged with parts of speech terms, and handed off to both a lexical analyzer and a semantic analyzer. Both analyzers generate questions separately and these questions are then combined and sent to the user to the right of Figure 3.

ConceptNet as Errol's Knowledge Base

Errol uses ConceptNet as its knowledge base (Speer et. al. 2017). ConceptNet is an open-source knowledge base that contains information on many everyday terms and represents them as concepts that are related to other concepts. For example, one of the ways a person may describe a doctor is as "someone who works at a hospital". ConceptNet encodes this information as *doctor atLocation hospital*, where *atLocation* is one of many relationships that may result from searching for *doctor*. Others include *capableOf* and *usedFor*, the latter of which says why you may go to a doctor. In addition to existing as a searchable web portal, ConceptNet also provides an API; Errol connects to ConceptNet through this API.

The business model on the input BMC may refer to any of the concept-relation pairs in ConceptNet. Thus, we encoded each concept-relation pair in Errol as a question. These questions make up the "secret sauce" that makes Errol work. Errol takes each noun found in the original BMC and sends an API call to ConceptNet. The JSON that is returned gives us all the concept-relation pairs with which we can generate a question.

Natural Language Analysis

Errol uses NLTK for preliminary processing of the input, namely the entries under VP and CS in the BMC (Bird et. al. 2009). In particular, Errol uses NLTK's part of speech (POS) tagging component to determine the appropriate parts of speech for the code to use. In order to determine what questions to ask, Errol performs lexical and semantic analysis upon its input, as illustrated in Figure 3.

Lexical Analysis

Errol's lexical analysis begins with a JSON object from the frontend that defines all the inputs and connections between BMC entries. These inputs are then tagged with a part-of-speech tag by NLTK's POS tagger, and each rule in the code goes through the inputs and tags to determine whether certain questions need to be asked of the user. If a question needs to be asked, it is added to a master JSON, which is what gets sent to the semantic analysis process.

The lexical system has a list of questions that were manually written and associated with specific triggers. In order to try and make the dialogue between the user and Errol seem a little more natural, each trigger had several questions associated with it, where each question asked the same thing but in several different ways. This way, the user would not necessarily get back the same exact question on every iteration, but the questions would revolve around the same topic from iteration to iteration if necessary.

Semantic Analysis

The semantic analysis process is similar to the lexical analysis process. NLTK's POS tagger is again used to tag the given inputs, but here, the nouns in particular are what are considered. For each noun, a query to ConceptNet is made and each of the relations that ConceptNet has on the noun are used to generate templated questions with the appropriate values plugged in. For example, if *counselor* is used as a search term to ConceptNet, then example relations would be *typeOf* and *capableOf*. Duplicate questions are pruned out and the master JSON is updated before questions are returned to the user.

Preliminary Results

Table 1 provides an example of the questions generated by Errol for the business model illustrated by BMC 709 in Figure 2. Errol's questions in Table 1 are organized by the CS and VP inputs that led to those questions. Goel et al. (2020) provides the results of running Errol on several business models.

For BMC 709 shown in Figure 2, we compared the questions Errol asked of the startup team with questions that human coaches had asked at the I-Corps cohort in Puerto Rico in 2018. Human coaches tend to ask in-depth and thought-provoking questions that induce reflection. They also challenge the ideas behind business models, whether it is the usefulness of what is written on the BMC as it relates to the business model, underlying assumptions and motivations, word choice, or even quantifiers. To facilitate being able to ask these kinds of questions, they ask for more detail, but only where things are too vague to understand what the intentions and goals of the startup team are. The main unifying theme is that the questions of human coaches are all backed by a strong understanding of the concept of *business*. While Errol has the same underlying aim as human coaches, its current iteration approaches the teaching challenge from a very different direction, particularly when it comes to the output of the semantic analysis. Unlike how human coaches try to encourage cutting back on things to highlight the more crucial parts of a business model, Errol's semantic analysis emphasizes questions that prompts the user to expand on ideas. The questions it asks are also more hollow "templated" questions that address the basic "who", "what", "where", "when", and "how" questions, but they do not really reflect any deep understanding of the BMC as a whole. We note that Errol may even output contradictory questions for the same BMC. Finally, the questions generated by lexical analysis are more similar to expert reviewer questions compared to those generated by semantic analysis.

Related Work

Insofar as we know, Errol is the first virtual coach that operates as a Socratic question-asking agent in the domain of entrepreneurship. Here, we briefly situate this work in research on four related topics: computational creativity, knowledge-based AI, intelligent tutoring systems, and question asking. In regard to computational creativity, the open-ended domain of entrepreneurship is similar to some respects to that of design: the problem in both cases is ill-defined, the evaluation criteria are difficult to operationalize, and the evaluation is delayed. Thus, in both design and entrepreneurship, the problem and the solution co-evolve: the problem specification gets revised as the solution takes form. The creative domains of entrepreneurship and design are similar in another respect: both engage the generation of models, business models in entrepreneurship and product models in design (Goel 2013). The major differences between them are that entrepreneurship requires a large investment upfront, the chances of success are low, and the cost of failure can be high. Thus, the margin for error is small.

In reference to knowledge-based AI, Errol makes extensive use of ConceptNet (Speer, Chin and Havasi 2017) that grew out of long-standing efforts to capture commonsense knowledge for interactive applications (Minsky 2004; Lieberman et al. 2004). It is designed to represent general knowledge involved in understanding language and it allows applications to better understand the meanings behind words. The knowledge graph consists of nodes that represent phrases and weighted edges that represent relations between two nodes. Entries in the knowledge graph include pointers to/from many external knowledge bases such as OpenCyc and WordNet. ConceptNet helps Errol understand a learner's input as it represents the relationships between the phrases. It can also allow Errol to advise on future directions by traversing other edges in the knowledge graph the student might not have explored.

Errol also represents AI research on education as well and can be viewed as an intelligent tutoring system. Graesser, Conley & Olney (2012) and VanLehn (2011) present two overviews of research on AI in education in general and intelligent tutoring systems in particular. However, unlike many intelligent tutoring systems that address well-defined closed-world problems with a single correct answer, Errol addresses an ill-defined, open-ended problem with no single correct answer known in advance. Further, while most tutoring systems support teacher-guided pedagogy in K-12 education, Errol is intended to support self-directed andragogy for life-long learning.

Question asking is receiving increasing attention in the literature on AI in education, knowledge-based AI, and computational creativity. (Kurdi et al. 2020) provide a recent review of automatic question generation, while Liu, Calvo & Rus (2014) describe the use of question asking to support critical thinking. Kearsley (1976) and Nielson et al. (2008) provide an early and a recent taxonomy of questions, respectively. Graessar's (2016) reflections on his AutoTutor project suggest that what questions an AI agent asks is more important than how it asks them; Errol takes a similar stance. In regard to question asking in computational creativity, our work is closest to that on RoboChair (Pollack et al. 2015), which uses templates to ask questions of scientific paper presentations by mimicking a human chair of a scientific conference.

In some ways, Errol may be considered as the opposite of Jill Watson (Goel and Polepeddi 2018), the first virtual teaching assistant. Jill answers learner's questions; Errol asks questions of learners. Jill replies to very specific questions with precise answers; Errol asks questions in response to ill-defined, vague, and evolving business models on a BMC. Indeed, Errol is more akin to the virtual research assistant VERA (An et al. 2020) than Jill Watson. VERA is an interactive open learning environment for inquiry-based modeling; it helps a learner construct conceptual models of natural phenomena, evaluate the models by simulation, and revise the conceptual and simulation models as needed. In VERA, a learner learns by doing and by reflection. Errol too is an interactive open learning environment for inquiry-based research, and it too helps a learner by doing and by reflection.

Discussion

Errol has several limitations compared to human coaches as indicated in the subsection on "Comparison of Errol's Question Asking with Human Coaches". Here, we present a brief critique of three aspects of Errol's information processing that also indicate directions for further work in the near future: the heavy reliance on NLTK's POS tagging component, the semantic analyzer's focus on nouns, and the lack of truly iterative question generation.

Errol's current architecture relies heavily upon the POS tagging component from the NLTK library. The lexical analyzer parses through the POS tags to return questions that the learner would need to reflect on to modify their entries in BMCs. Hence, the accuracy of the questions returned from lexical analysis is heavily correlated to the accuracy of the POS tags. The semantic analyzer also uses the subjects and nouns from the POS tags to query ConceptNet and search relevant terms and relationships. Given that many novice entrepreneurs do not write grammatically correct phrases or sentences, accuracy of the POS tags may be called into doubt. Improving the accuracy of the POS tags can be difficult as NLTK is an external library. Thus, a useful future action would be to develop a method of accurately parsing BMCs that contain grammatically incorrect phrases.

While it has been advantageous for the semantic analyzer to focus initially on nouns, a truly robust question generation system needs to be able to understand every input statement in its entirety. The focus on nouns allowed for the rapid development of initial algorithms for question generation. Yet, other parts of speech can influence what questions should be asked. For example, there may be a difference between asking questions about "small profit margins" versus "large profit margins". Thus, future iterations of Errol should incorporate techniques like verb and adjective recognition to enhance the set of questions that are produced and asked of the learner.

Truly iterative question generation may be characterized as a technique where a prior set of questions that Errol asks has influence upon a subsequent set of questions. In an iterative system of collaborative argumentative dialogue, if Errol determines that it should pose the same question several times, it likely needs the ability to choose prioritizing that question the next time or to ask fewer questions in total to make that one question stand out. The lack of calling attention to a particularly significant question makes it difficult for the user to see if they are making progress. A future version of Errol will need to address this shortcoming.

Conclusions

On the spectrum of creativity, entrepreneurship can be thought of as a new frontier. As entrepreneurship continues to grow in its allure, so too does the desire to learn how to design business models and the need for expert coaching of such designs. We see within this the need for virtual coaches to augment the coaching process. In this paper, we presented Errol as a preliminary startup virtual coach that seeks to simulate the experience of receiving feedback from an expert human coach. Would-be entrepreneurs tend to make the same kinds of mistakes in the various stages of developing their business models. Human coaches ask questions in the Socratic fashion to help the learners reflect on their mistakes and improve the business models. Errol similarly makes use of typical errors in developing business models to ask questions about a given business model and to help the learner learn by reflection on the feedback. Errol uses both lexical and semantic analysis to generate relevant questions, and in doing so, acts like a sounding board for the would-be entrepreneur. We intend to introduce Errol in a class on entrepreneurship and to assess the students' use of the virtual coach.

Acknowledgements

Several collaborators contributed to the Errol project over the last five years including Arup Arcalgud, Siddharth Gulati, James Howe, Nikhita Karnati, Jongmin Lee, Aaron Mardis, and Jae Ro. We thank the US NSF I-Corps for its support for this work. KM, the last author on this paper, is a lead instructor with the I-Corps program. We have named our virtual coach after Errol Arkilic, the first I-Corps program director. This research has been supported by internal seed grants from Georgia Tech. This paper is an abridged and much revised version of Goel et al. 2020.

References

An, S., Bates, R., Hammock, J. Rugaber, S., Weigel, E. & Goel, A. (2020) Scientific Modeling Using Large Scale Knowledge. *In Procs. 21st International Conference on AI in Education (AIED'2020)*, pp. 20-24

Besold, T., Schorlemmer, M., & Smaill, A. (2015) *Computational Creativity Research: Towards Creative Machines.* Atlantis Press.

Bird, S., Loper, E., & Klein, E. (2009), *Natural Language Processing with Python*. O'Reilly Media Inc.

Boden, M. (1990). *The creative mind: myths and mechanisms.* London: Weidenfeld and Nicholson.

Carland, J., & Carland, J. (2004). Economic development: Changing the policy to support entrepreneurship. *Academy of Entrepreneurship Journal*, 10, 105–114.

Candy, L., & Edmonds, E. (2002) *Explorations in Art and Technology*. Springer-Verlag, London (2002).

Cohen, H. (1995) The Further Exploits of Aaron, Painter. *Stanford Humanities Review* 4.2.

Cope, D. (2005) *Computer Models of Musical Creativity*. MIT Press, Cambridge.

Fitzgerald, T., Goel, A., & Thomaz, A. (2017) Human-robot co-creativity: Task transfer on a spectrum of similarity. In Procs. *Eighth International Conference on Computational Creativity (ICCC)*.

Forest, C., Sivakumar, R., Vito, R., Saxena, R., Harris, J., Perkins, R., Davidson, D., Ramachandran, K., McGreggor, K., Olufisayo, O., Klaus, C., & McLaughlin, S. (2021) CREATE-X: Toward student entrepreneurial confidence. *IEEE Potentials*, 40 (3), 14-22.

Goel, A. (2013) One Thirty Year Long Case Study; Fifteen Principles: Implications of an AI Methodology for Functional Modeling. *AIEDAM* 27(3): 203-215.

Goel, A. & Chandrasekaran, B. (1992) Case-Based Design: A Task Analysis. In *Artificial Intelligence Approaches to Engineering Design, Volume II: Innovative Design*, C. Tong and D. Sriram (editors), pp. 165-184, San Diego: Academic Press.

Goel, A., Hong, S., Kuthalam, M., … & McGreggor, K. (2020). Errol: A virtual coach for question asking and enabling learning by reflection in startup engineering. *Georgia Institute of Technology Technical Report.* (https://smartech.gatech.edu/handle/1853/63976).

Goel, A., & Polepeddi, L. (2018) Jill Watson: A virtual teaching assistant for online education. In Dede, C., Richards, J., and Saxberg, B., (Editors) *Education at scale: Engineering online teaching and learning.* NY: Routledge.

Graesser, A. (2016) Conversations with AutoTutor Help Students Learn. *International Journal of Artificial Intelligence in Education*, 26(1): 124 -132. 2016.

Graesser, A., Conley, M., & Olney, A. (2012) Intelligent tutoring systems. In: Harris, K., Graham, S., Urdan, T. (Eds.), *The APA Educational Psychology Handbook.* American Psychological Association, Washington, DC.

Huang-Saad, A., Fay, J., & Sheridan, L. (2017). Closing the divide: Accelerating technology commercialization by catalyzing the university entrepreneurial ecosystem with I-corps. *Journal of Technology Transfer*, *42*, 1466–1486

Kearsley, G. (1976) Questions and Question Asking in Verbal Discourse: A Cross-Disciplinary Review. J*ournal of psycholinguistic research*, vol. 5(4), pp. 355 -375, 1976.

Kuratko, D. (2005). The emergence of entrepreneurship education: Development, trends, and challenges. Entrepreneurship Theory and Practice,29(5), 577–598.

Kurdi, G., Leo, J., Parsia, B., Sattler, U., & Al-Emari, S. (2020). A systematic review of automatic question generation for educational purposes. International *Journal of Artificial Intelligence in Education* 30(1):121–204.

Lieberman, H., Liu, H., Singh, P., & Barry, B. (2004) Beating commonsense into interactive applications. *AI Magazine* 25(4): 63-76.

Liu, M., Calvo, R, & Rus, V. (2014). Automatic generation and ranking of questions for critical review. *Journal of Educational Technology & Society* 17(2):333–346.

Minsky, M. (2000) Commonsense-based interfaces. *CACM* 43(8): 67-73.

Nielsen, R., Buckingham, J., Knoll, G., Marsh, B., & Palen, L. (2008) A Taxonomy of Questions for Question Generation. In Procs. Workshop on The Question Generation Shared Task and Evaluation Challenge, Arlington, VA, USA, pp. 318-320.

Osterwalder, A., & Pigneur, Y. (2010) *Business model generation: A handbook for visionaries, game changers, and challengers*. John Wiley & Sons.

Pearce, M. & Wiggins, G. (2014) Expectation in melody: The influence of context and learning *Music Perception* 23 (5), 377-405

Pérez y Pérez, R., & Sharples, M. (2001) *MEXICA: A computer model of a cognitive account of creative writing. Journal of Experimental and Theoretical Artificial Intelligence*, 13119-139

Paul, R., & Elder, L. (2007). *The thinker's guide to the art of Socratic questioning*. Tomales, CA: The Foundation for Critical Thinking Press.

Pollak, S., Lesjak, B., Kranjc, J., Podpecan, V., Lavrac, N. et al. (2015). RoboChair: Creative assistant for question generation and ranking. In Procs. *IEEE Symposium Series on Computational Intelligence*, 1468–1475. IEEE.

Schank, R., & Cleary, C. (1995) Engines for Education. Routledge. Available as an ebook at https://www.engines4ed.org/hyperbook/

Speer, R., Chin, J., & Havasi, C. (2017) ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *Procs. Thirty-First AAAI Conference on Artificial Intelligence (AAAI 2017)*, pp. 4444 -4451.

Sternberg, R. (Editor, 1999) *Handbook of Creativity*. Cambridge University Press.

VanLehn, K. (2011). The relative effectiveness of human tutoring, Intelligent Tutoring Systems, and other tutoring systems. *Educational Psychologist*, 46, 197–221.

Veale, T. (2012). *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity.* Bloomsbury Academic, London.

Veale, T., & Cardoso, A. (Editors, 2020) *Computational Creativity. The Philosophy and Engineering of Autonomously Creative Systems.* Springer.