Exploring CC in XR: Visualizing Creative Conversation Topics to Facilitate Meaningful Face-to-Face Interaction

Hunter Harris, Makayla Thompson, Isaac Griffith, Paul Bodily
Department of Computer Science
Idaho State University
Pocatello, Idaho 83209 USA
harrhunt@isu.edu, thommaka@isu.edu, grifisaa@isu.edu, bodipaul@isu.edu

Abstract

With its ability to capture, model, and generate real-and-virtual combined environments, Extended Reality (XR) opens unique opportunities for designing creative agents. In this paper, we present HeyLo, a system for generating meaningful, visual conversation cues in XR. HeyLo autonomously analyzes user tweets to identify common interests and visualizes these interests as conversation cues using emoji. We argue that this system demonstrates creativity using four metrics that characterize novelty, value, and intentionality in the domain of conversation cues: specificity, inter- and intra-user variance, and relevance. We discuss several potential research questions that we hope to answer in the future using this system and the broader implications of a creative system that is capable of bridging arbitrary interests to innovate its own creative ideas.
Source: `github.com/harrhunt/HeyLo`

Introduction

The field of Computational Creativity (CC) devotes itself to "the art, science, philosophy and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative" (Wiggins 2006). Much great work has been done to develop *theory* of CC, examining general questions such as: "*What is creativity?*" (Boden 2003); "*What are computational frameworks and models for creativity?*" (Wiggins 2006); "*How do we measure creativity?*" (Ritchie 2007); and "*How do creative systems achieve autonomy?*" (Jennings 2010). In tandem with these contributions, the field has fueled innovation of *applied* CC systems in domains as varied as culinary arts (Morris et al. 2012); linguistic constructs (Veale 2018); visual art (Colton 2012); poetry (Toivanen et al. 2012); narrative and story telling (Pérez y Pérez and Sharples 2004); mathematics (Pease, Guhe, and Smaill 2010); software engineering (Colton, Powley, and Cook 2018); and Rube Goldberg machines (Xiou Ge and Varshney 2018). An important symbiotic relationship exists between these theoretical and applied contributions: the theory informs the application, and the applications inspire new inquiry and discussion of theory.

Because application is essential for the evolution of new theory, emphasis is rightly placed on encouraging research across and between an ever-widening spectrum of creative domains (Loughran and O'Neill 2017). It is one of the exciting challenges of the emergent field of CC to identify novel, valuable, and untapped domains in which to apply itself.

While there exist long-standing domains of creativity into which CC has yet to make an entrance, it occasionally happens that technology introduces an entirely new *medium* for creativity. This allows existing creative domains to take on new forms and provides computational systems novel access to new forms of creativity. An example of this is seen in the advent of extended reality (XR), an umbrella term that refers to all real-and-virtual combined environments (e.g., AR, MR, VR). Although XR dates back to the 1960s, commercially available and affordable XR to private consumers is a relatively recent phenomenon. This has prompted a significant increase in the demand for XR content (Moore 2017). In addition it has also opened new avenues of research and industry, including for example the use of XR to address issues of social isolation in older adults (Lin et al. 2018). Traditional production techniques have failed to meet this demand, leading to complaints about the lack of high-quality XR content. In response, some researchers have suggested that the solution lies in the development of procedural content generation (Tree and Malizia 2019).

The unmet demand for novel, creative XR content suggests unique and expansive opportunities for the field of computational creativity: opportunities for novel applications of CC, for evaluating how users interact with CC, and for identifying new lines of inquiry in CC theory. Although XR provides new opportunities for CC in familiar domains such as music, narrative, and visual art, of greater interest to the field is the question of "What forms of creativity have historically been inaccessible to computers to which XR provides unique access?" To explore this question, we set out to design a CC system for XR designed to leverage the unique purposes and strengths of this novel medium. The XR medium creates a space in which virtual and real worlds overlap to enhance real world interactions. It can capture, model, and influence interactions between people and organizations. How do these interactions relate to creativity?

Social psychology research demonstrates that face-to-face interaction produces up to 34 times more effective interactions than digital communication (Roghanizad and Bohns 2017). Research also demonstrates that a significant role in

Proceedings of the 11th International Conference on Computational Creativity (ICCC'20)
ISBN: 978-989-54160-2-8

the success of these interactions (from the viewpoint of an engaging party) is the degree of perceived similarity with other entities (Hampton, Fisher Boyd, and Sprecher 2019). Intuitively, this makes sense. We recognize the value of networking, building relationships of trust, and establishing common ground to improve the success of our interactions. Creativity is required to identify ways to establish common ground with others, and humans are not all equal at knowing *what* to say and *how* to say it. The "art of conversation" might essentially be defined as possessing the ability to engage in and maintain conversation about topics that are novel, valuable, and intentional—attributes that have been repeatedly used to define creativity (e.g., see (Ventura 2016; Wiggins 2006)). Can we conceive of a CC system that possesses this same ability?

We choose to focus on the simple task of generating meaningful conversation cues that effectively engage two parties in meaningful conversation. While this task can be done in mediums outside of XR (e.g., on social media), the task in the XR medium is unique: whereas on social media users largely interact with those who they intentionally seek out, in the real world people more frequently interact with those who they have *not* sought out or perhaps those who they might have otherwise intentionally avoided. The XR medium allows us to ask the question "Can a CC system be designed to suggest novel, valuable, and intentional conversation cues to engage two arbitrary parties in conversation that is meaningful to both *parties*?" We intentionally specify *parties* because the task we are describing is not unique to interactions between individuals. Whether it is two people meeting, two businesses interacting, a potential customer walking by a retail store, or a virtual chatbot, the problem is the same. Previous work has addressed this problem, but under the added assumption that interests are predefined and explicitly available (Nguyen et al. 2015; Jarusriboonchai et al. 2015). In our work, we consider the automated identification of interests to be an essential part of the problem to be solved. Our work is also unique in the visualization of interests using emoji.

In this paper we describe an extended reality computational creativity (XRCC) system, *HeyLo*, that operates in XR to identify potential topics of conversation between an XR user and another person encountered by the user in the XR medium and then overlaying visual representations of these interests using labeled emoji. To evaluate the creativity of the system, we define measures of novelty, value, and intentionality for artefacts generated by the system and apply these measures to comparatively evaluate the creativity of several different versions of HeyLo. We demonstrate examples of artefacts generated by HeyLo for real-world users. We discuss our future research agenda into XRCC and the implications of CC of a system that, more than merely finding common interests, bridges seemingly-incompatible interests to propose novel concepts.

Methods

In this section we describe the design and operation of the HeyLo XRCC system. We first provide a high-level overview of the system. At the heart of this system is a model which attempts to identify interests from a set of user tweets that when used as topics for conversation evoke a sense of novelty, value, and intentionality. For purposes of illustrating different levels of creativity in solving this problem, we outline five different approaches for identifying interests. Finally we define four metrics by which we comparatively evaluate the performance of each of these five approaches.

HeyLo System Overview

The HeyLo system runs on an MR[1] headset worn by a user with cloud support. Taking as input the image of a person from the headset camera, the system identifies a second user and then computes a set of weighted keywords representing shared interests between the two users. Each interest is paired with a representative emoji that the system overlays on the headset screen (see Figure 1).

To explicate where the creativity lies in this system and describe in further detail its implementation, we analyze HeyLo in terms of the FACE model (Colton, Pease, and Charnley 2011). In the FACE model, the creative behavior of a system is defined as a tuple of generative acts, containing 0 or 1 of eight possible types of creativity. Based on the behavior described in the system overview, we argue that the tuple for the HeyLo system is described by the tuple $< F^g, A^g, C^g, E^p, E^g >$. We consider the elements of this tuple in an order which best helps to describe the system and its creativity.

The Concept, C^g The concept of the artefact generated by the HeyLo system is a set of visual conversation cues. A visual conversation cue is defined as a text label representing a common interest between two users and an image (e.g., emoji) that represents the common interest. It is a set of cues that represents a concept because creativity is also required to ensure that cues within a set relate appropriately with each other (e.g., non-redundant, diverse, etc.).

The Expression of the Concept, E^g Figure 1 illustrates an expression of the concept expressed by HeyLo: a halo of labeled emoji around the face of a person. Of particular importance for the expression of visual conversation cues in HeyLo is that this expression occurs in the MR medium which allows the expression to take inspiration from and merge with the user's normal field of vision. Expressing the concept in the MR medium allows the system to better achieve its intention: help the MR system user to identify effective means for starting a conversation in an arbitrary encounter without becoming a distraction to either user[2].

The Method for Generating Expressions of a Concept, E^p HeyLo's method for generating visual conversation cues hinges on generative models for three fundamentally creative tasks: first, identifying appropriate interests for a

[1]Although HeyLo is specifically designed for MR (which gives the user access to their normal field of vision), the system can just as easily be used with AR or VR.

[2]Current MR headsets would most certainly constitute an explicit distraction. With this statement, we are envisioning future MR devices that are as inconspicuous as glasses or contact lenses.

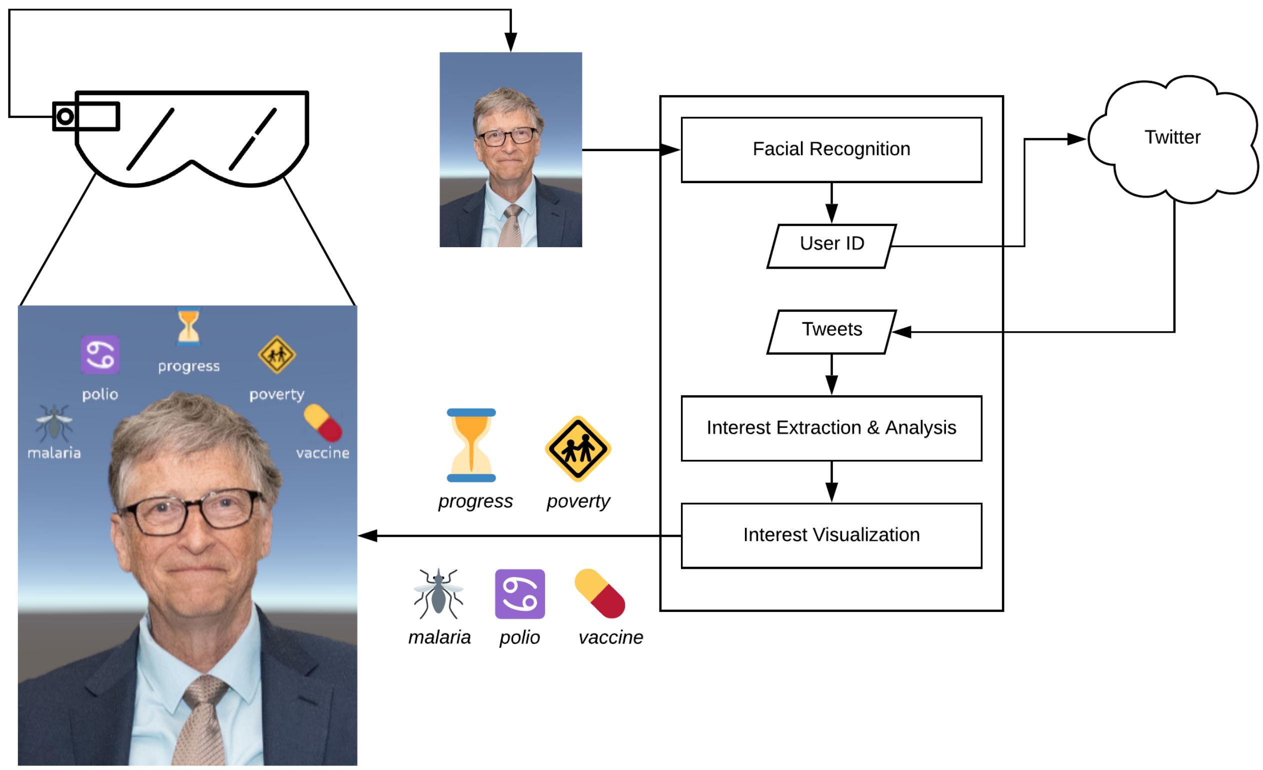

Figure 1: An overview of the HeyLo system. Given an image of a person (e.g., from an XR headset), the system recognizes the user in the image, gathers tweets and extracts interests for both the image user and the headset user, and outputs visualized conversation cues based on common interests between the two users. Picture of Bill Gates is in the public domain.

single user; second, identifying an appropriate method for finding common interests between users; and third, identifying appropriate images for visualizing the common interests that it generates. We use five different interest identification models and compare the results from each to find the best model for identifying a single user's interests. We will denote the interest identification model with M. The process below is the same for each model M we use.

Details of the method which HeyLo uses to generate visual conversation cues are as follows. Given the user u_{me} wearing the MR headset and an input image f of a person taken from the headset camera,

1. Identify from the set of all users U (i.e., all available social media users) the user u represented in f using facial recognition[3].
2. Compute a set of weighted interests for u as follows:
   (a) Retrieve 1000 of the most recent social media posts created by u and represent the content of these posts as a single multiset of words S. Filter S for stopwords, URLs, and non-alphabetic characters, and for each word $s \in S$, replace s with the NLTK (Bird, Bird, and Loper 2016) WordNetLemmatizer lemmatized form of s.
   (b) Apply interest identification model M (described below) on S to obtain a set of pairs $I_{M,u} = \{(k_1, w_1), \ldots, (k_i, w_i), \ldots, (k_n, w_n)\}$ where k_i is a keyword or interest and $w_i \in \mathbb{R}_{\geq 0}$ represents a weight or level of interest for k_i. Normalize weights for interests in $I_{M,u}$ to range from 0 to 1 (i.e., divide each weight by largest weight in $I_{M,u}$).
3. Repeat step 2 to create $I_{M,u_{me}}$ for the MR headset wearer u_{me}.
4. Compute a set of shared, weighted interests $I_{M,\{u_{me},u\}}$ from $I_{M,u_{me}}$ and $I_{M,u}$ such that $I_{M,\{u_{me},u\}}$ is the set of all pairs (k_i, w_i) where k_i appears as an interest in both $I_{M,u_{me}}$ and $I_{M,u}$, and w_i is the product of the two weights for k_i in $I_{M,u_{me}}$ and $I_{M,u}$.
5. Reduce $I_{M,\{u_{me},u\}}$ to the pairs (k_i, w_i) with the l highest weights w_i.
6. For $(k_i, w_i) \in I_{M,\{u_{me},u\}}$ select an emoji $e_i \in E$ (where E is the set of all emoji) as follows:
   (a) Use word2vec (Mikolov et al. 2013) to compute a vector representation v_i for interest k_i.

[3]This step is envisioned as part of future work and is not currently employed as part of HeyLo.

Proceedings of the 11th International
Conference on Computational Creativity (ICCC'20)
ISBN: 978-989-54160-2-8

(b) Identify an emoji e_i from v_i by computing

$$e_i = \operatorname*{argmin}_{e \in E} d(v_i, v_e)$$

where v_e is the vector obtained by applying word2vec to the text label for an emoji e (where the label is multiple words, v_e is the average of the vectors for each word), and d is the function computing the cosine distance of two vectors (i.e., a measure of dissimilarity).

(c) Add the triple (k_i, w_i, e_i) to the return set R.

7. Display the emoji e_i with label k_i for each triple $(k_i, w_i, e_i) \in R$ to form a halo around the face of u in f (e.g., see Figure 1).

An Aesthetic Measure, A^G HeyLo possesses an aesthetic that favors interests that result from the overlap of the interests of two users. Furthermore, of the set of common interests, the system prefers those interests which are more heavily weighted by each user independently (as determined by the interest-extraction method). This is a temporary aesthetic until the system's functionality has been expanded to be able to bridge non-overlapping interests. We plan to also incorporate the evaluation metrics of specificity, relevance, and inter- and intra-user variance (defined below) as an explicit part of the system's aesthetic in future iterations.

An Item of Framing Information, F^g HeyLo possesses a highly decomposed conceptualization model, which makes it easier for the system to describe its intentions and thinking in creating particular sets of conversation cues. It is not appropriate to provide this framing information in the MR medium (such would cause a distraction to the user), but providing the means by which the user can access this information is part of the fully-envisioned HeyLo system.

Evaluating Creativity in Conversation Cues

In developing HeyLo as a CC system, we discovered that several of the approaches we tried for identifying potentially creative conversation cues exhibited a lack of creativity. This was manifest in some of the approaches generating cues that failed the test of novelty, of value, or of intentionality. We found that this originated in decisions the system made for selecting an individual user's interests. To objectively analyze how well-suited a particular approach is for identifying individual interests, we developed four metrics that collectively capture the notions of novelty, value, and intentionality for a set of conversation cues derived solely from the interests of a single user. These metrics are: *specificity* (value), *intra-user variance* (novelty), *inter-user variance* (novelty), and *relevance* (intentionality).

Specificity Finding common ground between two people is easy if that ground is poorly specified, however, such generality in finding common interests is unlikely to create a shared perception of similarity between users. To be valuable, a conversation cue must be specific. How can we quantify specificity? Consider that for two words v and w, v IsA w indicates that v is more specific than w (e.g., "field lacrosse" IsA "sport"). We define $in\text{–}degree(k)$ for a keyword k as the number of words v such that v IsA k is a valid relationship catalogued in ConceptNet. From this we define the specificity of an interest k as

$$specificity(k) = 1/(in\text{–}degree(k) + 1)$$

Note that if k has an in-degree of 0 (e.g., as with k = "field lacrosse"), k cannot be further categorized or specified. In this scenario, k would receive the maximum specificity score of 1.0.

From this we can define the specificity of an interest identification model M. Let I_M represent the set of unique interests extracted by M across all users. Then the *specificity* of M is the average of the specificity values for each unique interest:

$$specificity(M) = \frac{\sum_{k \in I_M} specificity(k)}{|I_M|}$$

Intra-user variance Considering that HeyLo generates a *set* of visual conversation cues, the creativity of the system depends as much on the diversity of cues in the set as it does on the cues themselves. In returning a set of interests representative of a user, a good model will extract a set of *diverse* interests. For a set of interests $I_{M,u}$ extracted by model M for user u, we define the *variance* of $I_{M,u}$ as

$$variance_{intra}(I_{M,u}) = \sum_{(k_i,w_i),(k_j,w_j) \in I_{M,u}} D_C(v_i, v_j)$$

where $D_C(v_i, v_j)$ is the cosine distance between two vectors v_i and v_j representing interests k_i and k_j. Using the definition of intra-user variance for a set of interests, we define the intra-user variance of a model M as the average of the intra-user variance values across all users:

$$variance_{intra}(M) = \frac{\sum_{u \in U} variance_{intra}(I_{M,u})}{|U|}$$

Inter-user variance In addition to extracting a set of diverse interests, a good model will also extract diverse sets *across* users or, in other words, avoid repeatedly extracting the same interests for multiple users. To measure this inter-user variance, we find the sum of unique words for each user interest set divided by the total unique words across all user sets:

$$variance_{inter}(M) = \frac{|\bigcup_{u \in U} I_{M,u}|}{\sum_{u \in U} |I_{M,u}|}$$

Relevance The success of an interest identification model depends on extracting interests that are not only specific and varied, but which also reflect the user's actual interests. This final metric may be the most important of all, but is also one of the most challenging aspects to measure. A model's predictions for a user's interests can only be accurately assessed by the user him/herself. We are planning to conduct such a study as future work.

Interest Identification Models

The development of the HeyLo system included several iterations of testing of different interest identification models

Metric	Empath	Retrained Empath	Raw Word Count	Bayesian	Chi-square
Specificity	0.008 ± 0.007	0.020 ± 0.076	0.092 ± 0.238	0.685 ± 0.374	0.275 ± 0.365
Intra-user variance	0.173 ± 0.057	0.141 ± 0.052	0.171 ± 0.072	0.115 ± 0.068	0.135 ± 0.083
Inter-user variance	0.046	0.090	0.230	0.999	0.828

Table 1: Comparison of creative attributes exhibited by five interest identification models

M to see which model performed best according to the evaluation metrics of specificity, variance, and relevance. Each implementation uses a different method to identify the set of weighted interests:

1. *Empath* (Fast, Chen, and Bernstein 2016) - Empath maps word lemmas to 200 predefined categories (representing potential interests) thereby assigning a weight equal to the number of hits for each category.
2. *Retrained Empath* - A custom version of Empath adding 1000 categories from Facebook page categories[4].
3. *Raw word counts* - Word lemmas represent potential interests and the count for each lemma represents the weight associated with the interest.
4. *Bayesian* - Word lemmas represent potential interests, however the probability of the user u given the lemma k_i is used as the weight for k_i. Distributions were trained using the last 1000 tweets from each of a set of 500 highly-followed users [5].
5. *Chi-square* - Word lemmas represent potential interests, and the weight for interest k_i is computed as the chi-square contribution of the occurrence of k_i for user u with expected counts derived from data collected in the same manner as in the Bayesian approach.

Comparative Results

Comparative results can be found at `https://www2.cose.isu.edu/~bodipaul/research/heylo/`. Our goal in comparing different interest identification models was to identify which model produces artefacts exhibiting the most creativity. To evaluate specificity and variance, we calculated and averaged results for each approach over 500 of the most followed, publicly available Twitter handles. The results of these calculations are shown in Table 1. To evaluate relevance, we preselected five Twitter handles for five widely-recognized celebrities. We selected the celebrities from varying occupations and backgrounds to avoid returning similar interests for each example. We chose to perform the analysis on these users on the basis that their interests are generally well-known and therefore the results could be more easily assessed for how well the system's intention of identifying user interests was achieved.

By the specificity and variance metrics the Bayesian model *looks* to have achieved a significant amount of novelty and (to some extent) value. Looking at Dave Ramsey's results, for example, the Bayesian model extracts words that are specific and varied (e.g., *belay*, *godly*, *backache*, *variable*, and *toolbox*). These cues represent topics that most users do not often discuss. They also represent topics that are specific enough to avoid a conversation that is too general to carry meaning. These results, however, have very low relevance. Knowing that Dave Ramsey is a businessman, author, and renowned financial advisor, the interests extracted by the Chi-square model are significantly more relevant (e.g., *money*, *advice*, *financial*, *debt*, and *millionaire*). Inter-user variance scores that are at either extreme are undesirable because at the lower extreme (i.e., topics overlap) they lead to repetitive topics, and in the upper extreme (i.e., topics do not overlap) they lead to irrelevant topics. This leads us to conclude that inter-user variance scores that are not at either extreme are acceptable values for a given model. Both the default and retrained Empath models have very low inter-user variance scores reflecting that these models extract many of the same words across several users (low novelty). The Bayesian model has a very high inter-user variance score that leads to unique (high novelty) but irrelevant (low intention) topics. The raw word count and chi-square models both have intermediate values meaning they have an acceptable amount of inter-user variance to not suffer from the pitfalls of being at either extreme.

The low specificity score of the default and retrained Empath models betrays that these models also extract words which are relatively non-descript (e.g., *play* and *party*) and therefore likely to lead to conversations of low value. The raw word count model also suffers from low specificity (e.g., *thank*, *thanks*, and *people*). The Bayesian model has more specific words for each user, however the nature of this model leads it to prefer words that given the data are uniquely used by a particular user, even if these words are not the user's interests. As the amount of data in the model increases that the Bayesian model will improve, but in its current state, this model suffers from over-specificity.

The results of the chi-square model exhibit keywords that are varied and specific but without being too specific (see Table 2). The NULL emoji for Nick Offerman's word *pawnee* (a fictional city in a TV series featuring Offerman) was a result of the word *pawnee* not being in the Google News word2vec model. The model is unable to find the closest associated emoji because a vector for the word could not be determined. Future work will seek to address this issue.

[4] `https://www.facebook.com/pages/category/`

[5] `https://socialblade.com/twitter/top/500/followers`

User	*Interests* & **Emoji Visualizations**				
@DaveRamsey	*money* **money-bag.png**	*advice* **warning.png**	*financial* **bank.png**	*debt* **credit-card.png**	*millionaire* **man-farmer.png**
@realDonaldTrump	*transcript* **memo.png**	*witch* **woman-fairy.png**	*impeachment* **file-cabinet.png**	*democrat* **man-office-worker.png**	*shifty* **wavy-dash.png**
@tonyhawk	*skate* **ice-skate.png**	*birdhouse* **snowman.png**	*hawk* **owl.png**	*demolition* **building-construction.png**	*vert* **skateboard.png**
@Nick_Offerman	*nick* **man-bouncing-ball.png**	*mirth* **cat-with-tears-of-joy.png**	*berry* **cherries.png**	*woodworker* **man-artist.png**	*pawnee* **NULL**
@BillGates	*malaria* **mosquito.png**	*polio* **cancer.png**	*progress* **hourglass-not-done.png**	*poverty* **children-crossing.png**	*vaccine* **pill.png**

Table 2: Visualized conversation cues generated by HeyLo based on individual user interests

HeyLo Results using Chi-square Model

Our study allowed us to conclude that as a combination of relevance, specificity, intra-, and inter-user variance, the chi-square model was the best approach for identifying interests that would lead to creative conversation cues. Using this approach, HeyLo successfully finds meaningful topics of interest for two users and then identifies potential conversation cues from the overlap in these interests (see Table 3).

An example of where the system generated a good conversation cue is with the interest *ukulele* for Bill Gates and Nick Offerman. After looking into this more, Bill Gates sang while Warren Buffett played the ukulele in a video from 2016. Nick Offerman wrote a song on the ukulele and has even made his own ukuleles. A meaningful conversation between these two users might result from sharing their own experiences with ukuleles even though they seem to have very little in common.

Other results suggest potential areas of improvement in HeyLo. Some of the results such as *great* and *thanks* would not spark any meaningful conversations between users. To increase the quality of the conversation cues generated, a new method for finding connections between users is necessary. One solution to this problem is to creatively bridge seemingly disparate interests (discussed below).

Discussion and Conclusion

In developing our vision of an XRCC system for visualizing conversation cues, we established several milestones. HeyLo, as presented in this paper, accomplishes the first of these milestones, which is to design an XRCC framework that autonomously elicits user interests and visualizes conversation cues based on common interests between users with the intention of facilitating meaningful interaction between them. The current system may yet benefit from a refinement of the metrics used to measure different interest identification models (e.g., leveraging work done by Joho and Sanderson (2007)). An outline of our next steps are outlined in the following research questions:

1. Can HeyLo effect meaningful conversation through the proposal of visual conversation cues based on distinct yet compatible interests (e.g., *Switzerland* and *chocolate*)?
2. Can HeyLo propose visual conversation cues through the creation of novel interests that form from bridging disparate interests (e.g., *computer science* and *public defense*)?
3. Can HeyLo's intention be augmented to account for the polarity (+/-) of a user's sentiment towards an interest?
4. Can HeyLo's intention be augmented to account for an individual's mood (e.g., based on facial expressions)?
5. Can HeyLo's creativity be used to suggest pairs of users who are likely to engage in meaningful conversation?

We will focus on questions 2 and 3 from the list above and discuss their significance to our system's expansion.

Bridging Seemingly Disparate Interests If common interests cannot be found, one solution is to redefine the space of user interests through the creation of novel user interests that would be common to both users. As an example, consider the following social encounter between one of the authors (a computer science professor) and a neighbor (a public defender). The meaningfulness of this particular interaction was initially stifled by an apparent lack of *common* interests between the author and the neighbor. The meaningfulness of the conversation dramatically increased, however, when the author began to seek for ideas for potentially bridging the two parties' disparate sets of interests and came up with the idea of teaching computer science to indigent defendants. This idea eventually became the impetus for the

	@realDonaldTrump	@tonyhawk	@Nick_Offerman	@BillGates
@DaveRamsey	*great* (3.46*e*-3) *insurance* (2.24*e*-3) *president* (1.80*e*-3) *money* (1.43*e*-3) *job* (5.91*e*-4)	*thanks* (1.96*e*-4) *workbook* (1.32*e*-4) *agent* (1.20*e*-4) *millionaire* (5.46*e*-5) *gleam* (5.45*e*-5)	*book* (4.78*e*-4) *berry* (3.06*e*-4) *preorder* (1.68*e*-4) *advice* (1.27*e*-4) *instruction* (5.83*e*-5)	*book* (2.82*e*-3) *advice* (1.36*e*-3) *teacher* (1.12*e*-3) *money* (7.82*e*-4) *save* (6.55*e*-4)
@realDonaldTrump	–	*call* (1.92*e*-4) *thanks* (1.40*e*-4) *unrelated* (6.47*e*-5) *finder* (4.71*e*-5) *love* (3.08*e*-5)	*fabrication* (6.65*e*-5) *salient* (5.14*e*-5) *book* (3.95*e*-5) *impeachment* (3.58*e*-5) *cherry* (3.41*e*-5)	*fight* (9.54*e*-4) *president* (4.44*e*-4) *energy* (4.17*e*-4) *book* (2.33*e*-4) *aid* (2.15*e*-4)
@tonyhawk	–	–	NULL *underhill* (2.57*e*-4) *sandpaper* (4.95*e*-5) *tony* (2.86*e*-5) *board* (9.33*e*-6) *proceeds* (8.11*e*-6)	*help* (3.31*e*-5) *world* (2.92*e*-5) *crum* (2.81*e*-5) *hawk* (2.66*e*-5) *neat* (2.12*e*-5)
@Nick_Offerman	–	–	–	*book* (2.88*e*-4) *world* (3.39*e*-5) *programmer* (3.37*e*-5) *ukulele* (2.93*e*-5) *optimist* (2.76*e*-5)

Table 3: Visualized conversation cues generated by HeyLo for pairs of users

creation of a new introductory computer science course in the local women's correctional center.

The process of taking known ideas and finding novel, valuable, and intentional bridges between them that lead us to new findings or applications of knowledge might be considered the essence of creativity itself. Our goal is to develop HeyLo into a system that explicitly models this process by bridging seemingly disparate interests in unique ways.

The ability to bridge seemingly disparate interests has significant applications for business and consumers. Bridging the dissimilarities between the consumer's interests and the business' offerings allows both the consumer and the business to be more efficient in their interactions. Consider as an example a user who is interested in computers and programming walks by a clothing store, but has no interest in clothing. The system could bridge the dissimilarity and suggest a pair of compression gloves that help with carpal tunnel syndrome. The user now has the opportunity to have a successful interaction with a business they otherwise would not have had. The challenge of bridging topics has been the subject of significant research (e.g., (Berthold 2012; Olsson et al. 2020; Nguyen et al. 2015)) which we plan to incorporate into future work.

Polarizing Interests We define a polarized interest as an interest together with an individual's sentiment toward that interest. Currently HeyLo disregards the polarity of a user's sentiment towards their identified interests. Incorporating polarity, however, gives the system an increased capacity for intentionality. There are many scenarios in which the system can use the polarity of interests to better achieve the intentions of the user. Consider the following examples:

- A user only wants to see common interests for which both users share the same polarity.
- A user wants polarity visualized so that they can be alerted to an individual's sentiment toward a particular topic (e.g., to approach controversial topics tactfully).
- A user wants to filter the interests they see based on polarity. This can go two ways: the user only wants to see interests towards which an individual feels favorably; or the user only wants to see interests towards which an individual feels unfavorably (e.g., for purposes of engaging with different points of view or for sparking debate).

From the examples above, there are many different intentions for social interaction that can be derived from using polarized interests. It is important to note that not all of the interactions described above are positive in their intentions, opening the need for ethical considerations in expanding HeyLo in this direction.

We have suggested that XR presents novel research opportunities for CC. As an example, we have presented HeyLo, an XRCC designed to autonomously generate visual conversation cues when encountering other users in XR environments. We have discussed measurements of specificity, variance, and relevance as means of evaluating novelty, value, and intentionality in this domain and demonstrated a comparative analysis of variations of the HeyLo system using these metrics. The system we have presented represents a basic framework for XRCC in which we hope to continue research into how to effectively find bridges between seemingly disparate interests in order to generate creative visual conversation cues.

References

Berthold, M. R. 2012. *Bisociative Knowledge Discovery: an introduction to concept, algorithms, tools, and applications*. Springer Nature.

Bird, S.; Bird, S.; and Loper, E. 2016. NLTK : The natural language toolkit NLTK : The Natural Language Toolkit. *Proceedings of the ACL-02 Workshop on Effective tools and methodologies for teaching natural language processing and computational linguistics-Volume 1*.

Boden, M. A. 2003. *The Creative Mind: Myths and Mechanisms, Second Edition*. Routledge.

Colton, S.; Pease, A.; and Charnley, J. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *Proceedings of the Second International Conference on Computational Creativity*, 90–95.

Colton, S.; Powley, E.; and Cook, M. 2018. Investigating and Automating the Creative Act of Software Engineering. *International Conference on Computational Creativity (ICCC)*.

Colton, S. 2012. The painting fool: Stories from building an automated painter. In *Computers and Creativity*.

Fast, E.; Chen, B.; and Bernstein, M. S. 2016. Empath: Understanding Topic Signals in Large-Scale Text. 4647–4657.

Hampton, A. J.; Fisher Boyd, A. N.; and Sprecher, S. 2019. You're like me and I like you: Mediators of the similarity–liking link assessed before and after a getting-acquainted social interaction. *Journal of Social and Personal Relationships*.

Jarusriboonchai, P.; Olsson, T.; Prabhu, V.; and Väänänen-Vainio-Mattila, K. 2015. Cuesense: A wearable proximity-aware display enhancing encounters. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, 2127–2132.

Jennings, K. E. 2010. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines* 20(4):489–501.

Joho, H., and Sanderson, M. 2007. Document frequency and term specificity. In *Proceedings of the Recherche d'Information Assistée par Ordinateur Conference (RIAO)*. Sheffield.

Lin, C. X.; Lee, C.; Lally, D.; and Coughlin, J. F. 2018. Impact of Virtual Reality (VR) experience on older adults' well-being. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*.

Loughran, R., and O'Neill, M. 2017. Application Domains Considered in Computational Creativity. In *Eighth International Conference on Computational Creativity (ICCC'17)*.

Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013 - Workshop Track Proceedings*.

Moore, T. P. 2017. What Steam's Data Reveals About The Health of VR's Ecosystem.

Morris, R. G.; Burton, S. H.; Bodily, P. M.; and Ventura, D. 2012. Soup Over Bean of Pure Joy : Culinary ruminations of an artificial chef. In *Proceedings of the Third International Conference on Computational Creativity*, 119–125.

Nguyen, T. T.; Nguyen, D. T.; Iqbal, S. T.; and Ofek, E. 2015. The known stranger: Supporting conversations between strangers with personalized topic suggestions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 555–564.

Olsson, T.; Jarusriboonchai, P.; Woźniak, P.; Paasovaara, S.; Väänänen, K.; and Lucero, A. 2020. Technologies for enhancing collocated social interaction: review of design solutions and approaches. *Computer Supported Cooperative Work (CSCW)* 29(1):29–83.

Pease, A.; Guhe, M.; and Smaill, A. 2010. Some aspects of analogical reasoning in mathematical creativity. In *Proceedings of the First International Conference on Computational Creativity*, 60–64.

Pérez y Pérez, R., and Sharples, M. 2004. Three computer-based models of storytelling: BRUTUS, MINSTREL and MEXICA. *Knowledge-Based Systems*.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17(1):67–99.

Roghanizad, M. M., and Bohns, V. K. 2017. Ask in person: You're less persuasive than you think over email. *Journal of Experimental Social Psychology*.

Toivanen, J. M.; Toivonen, H.; Valitutti, A.; and Gross, O. 2012. Corpus-based generation of content and form in poetry. In *Proceedings of the 3rd International Conference on Computational Creativity, ICCC 2012*.

Tree, D. J., and Malizia, A. 2019. Author-Driven Approaches to Computational Narrative Design for Games. In *Proceedings of the International Conference on Human-Computer Interaction*, 571–584. Springer.

Veale, T. 2018. A Massive Sarcastic Robot: What a Great Idea! Two Approaches to the Computational Generation of Irony. In *Proceedings of the Ninth International Conference on Computational Creativity, ICCC 2018*.

Ventura, D. 2016. Mere Generation: Essential Barometer or Dated Concept. *Proceedings of the Seventh International Conference on Computational Creativity, ICCC* (June):22–29.

Wiggins, G. A. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.

Xiou Ge, J. X., and Varshney, L. 2018. Computational Creativity for Valid Rube Goldberg Machines. *Proceedings of the Ninth International Conference on Computational Creativity (ICCC'18)*.