Contextual Bandits for Design: A Human-Computer Collaboration Approach

Janin Koch
Aalto University
Helsinki, Finland
Janin.Koch@aalto.fi

Abstract

Human-Computer Collaboration offers great potential for explorative and creative problem-solving strategies. While previous work in HCI and ML mainly focuses on exploiting either human or machine capabilities, the concept of collaboration suggests work on equal terms to achieve synergy effects. The uncertain nature of creative problems raises new questions regarding the adaptability of systems to changing objectives in iterative processes. We present a collaborative system for mood board design based on a state-of-the-art contextual bandit structure that is able to iteratively adapt to changing behaviors, and moves autonomously through solution spaces to propose suitable contributions. Besides the technical implementation, we discuss the need for further research on collaborative interaction behaviors between humans and machines.

Author Keywords

Collaborative AI; Design Inspiration; Interaction Concepts

ACM Classification Keywords

I.2.6. Computing Methodologies: Learning

Introduction

The idea of a machine that collaborates with humans has fired the imagination of computer scientists and engineers for decades. As early as 1960 J.R. Licklider wrote about machines and humans operating on equal footing and being able to 'perform intellectual operations much more effectively than a man alone'[1]. However, these benefits were hardly realized by the existing paradigms stemming from human-computer interaction (HCI) and machine learning (ML) research. In previous HCI work, systems and technology were often utilized as a means to support human-driven tasks where the system plays the role of an observant, serving instance within the interaction. In contrast, more technical approaches such as ML and Artificial Intelligence (AI) often exploit the potentials of technology with limited help from the user. Here the human plays a serving role in form of a feedback and evaluation instance in this setup. While these approaches extend the potentials of each agent, human and machine, these one-sided enhancements do not spark the potential synergy effects created through collaboration. What we envision are systems that pro-actively collaborate with humans on comparable terms.

Human-Computer Collaboration

Collaboration is 'a process through which parties who see different aspects of a problem can constructively explore their differences and search for solutions that go beyond their own limited vision of what is possible'[2]. Observations in fields like business, research and design showed significant improvements of problem-solving behaviours through the combination of skills [3–5]. Especially in innovative and creative processes does the diversity of skills improve the explorative finding of problem and solution definitions. The capabilities of current systems do reflect a huge potential in such processes by offering analytical and predictive power based on large scale data analysis. However, the uncertain and exploratory nature of these problems, like in design, where neither the final goal nor the complete design space is specified beforehand, requires the creation and rejection of potential solutions in an iterative manner [6].

Most AI systems today, however, require the definition of goal and environmental parameters in order to operate in an optimal way. The notion of Human-Computer Collaboration (HCC) therefore raises new questions and challenges for computational agents working in such underspecified adaptive environments. How can system objectives *evolve* within an iterative process instead of being predefined? How can systems adapt explorative and exploitative strategies to find novel solutions? While these questions mainly target the technical representation of creative processes, there is also a need to expand our knowledge of interactive computational behaviour. This includes abilities such as expressibility (express and understand reasoning), value alignment (diversify and narrow down the idea spaces), and agency (act on one's own behalf), and how these impact, mediate and facilitate HCC from human and system perspectives.

Cooperative Contextual Bandits for creating Mood boards

In our current work we explore the possibilities for gathering inspirational material together with an algorithm. In the following we will present our project, a computational structure capable of active learning under changing objectives, and highlight the future needs we envision in this direction.

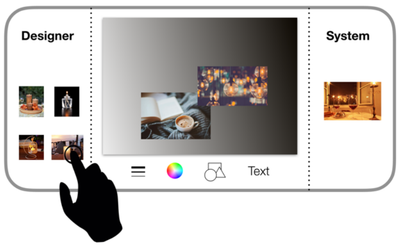

Figure 1: Sketch of the user interface

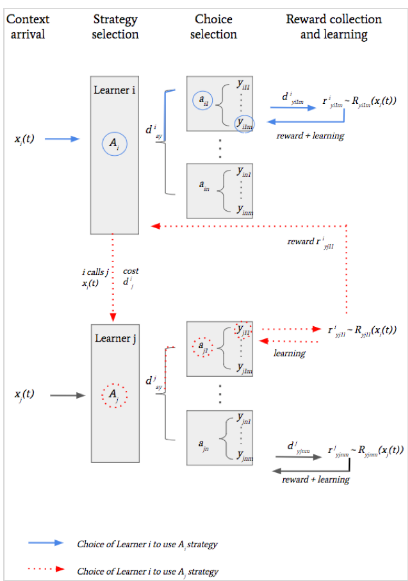

Figure 2: Structure of the Hierarchical Cooperative Contextual Bandit

Mood Board Design

Within the framing phase of a design process, requirements and inspirations are gathered to define possible characteristics of the later design. Especially collecting inspirational material helps designers to explore and narrow down the understanding of a given uncertain task. One common method in practice is the use of Mood Boards, a construction of visual references conveying the mood and impression that the final design should reflect [7]. Creating such mood boards is often a collaborative process where ideas are iteratively explored and rejected. The creative space considered depends on the individual contributors as well as their shared understanding of the task. When we abstract this task, we can identify three main aspects within this process: (a) the cognitive, creative and empathic understanding of the aimed goal and design, (b) the availability of material that reflects this understanding and (c) the iterative, creative process of exploring the potential idea space and narrowing it down based on a common understanding.

In collaboration, each agent (here human and system) tries to exploit its strengths to contribute in the best possible way to the final results. While the empathetic understanding of goals and impressions is a skill professional designers are trained in, the choice and availability of material is dependent on the creativity of the agents as well as on the means to retrieve them. A system-agent's ability to analyze large amount of data could contribute to this aspect. Further, all agents within the collaboration must be able to iteratively explore and narrow down the creative space. In our project, a designer and a collaborating system would work at one mood board by iteratively adding and rejecting images as shown in Fig 1.

Implementation

We assume that the purpose of the mood board is known to both agents, in an abstract format like: *"The design of an energetic and modern website for a new car brand".* Based on this information both agents start looking for images related to this topic as a starting point for the current mood board. The designer will have common tools like image search, color, text and shape tools for adding elements on a shared canvas. The system-agent applies an association algorithm to explore topics related to the given purpose, and can also perform large-scale image search. While experience and empathy guide the designer's decision to add or remove an image, the system-agent has to be able to make a comparable decision.

To fulfill these requirements and adapt to changing objectives, a collaborative system needs to (1) make context-dependent decisions, (2) change its explorative strategies, (3) adapt to its current collaborators. For this purpose we extended the state of the art Cooperative Contextual Bandit system presented by Tekin and van der Schaar [8], to reflect the structural necessities of the current task. The resulting structure of our Online Learning Hierarchical Cooperative Contextual Bandit to address the aforementioned three features is presented in Fig. 2.

The context is defined in visual dimensions of the current mood board canvas *(MB)*. We assume for now that the designer provides the first image of a context x_n, $MB = \sum x_n$. We consider the visual features *dominant Color, Saturation, color Temperature* and *Texture* as well as *dominant color Distance* for describing the context. $MB = \langle C_{15}, S_3, Temp_3, T_3, D_5 \rangle$

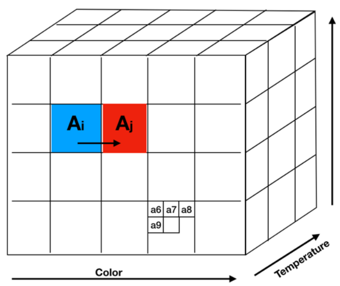

Figure 3: Simplification of the creative space and moving from strategy A_i to A_j

This state-space *MB* is uniformly sliced into hypercubes, $A_n = \langle C_5, S_2, Temp_2, T_2, D_3 \rangle$ and each cube A_n represents a certain type of mood board design, also called design strategy (e.g. A_i in Fig. 2). Each strategy A_n has a set of subagents $a_{n1}, \dots, a_{nm}$ representing subsets of the considered visual features, e.g. $a_{11} = \langle C_{15}, Temp_3 \rangle$, with learned probability distribution y_{11}. When a certain context is observed, the current strategy A_n has to decide whether to exploit one of its own subagents (e.g. a_{n1}) or to explore the creative space by referring the decision to another strategy (e.g. A_j shown in Fig. 3). This allows the system to change suggestion strategies during the iterative process, depending on the current context of the mood board.

Generally said, the algorithm chooses the subagent a_{nm} that presents the highest probability to be the optimal one. In order to choose one of these strategy options, we apply a Thompson sampling algorithm, which offers a better balance of exploration/exploitation than previous techniques and whose total payoff has been experimentally shown to be close to the optimal strategy [9]. When the system-agent suggests an image for the mood board, the human designer can either accept or reject it, resulting in a Bernoulli reward function of the form $R = \langle 0,1 \rangle$. Following previous work[9], we use a corresponding Beta-prior $Beta(\alpha, \beta)$ with parameters $\alpha > 0,\ \beta > 0$ to model the strategy selection process whose parameters are updated for every distribution $y_{n1}, \dots, y_{nm}$ within A_n at every step t with $r(t)$, $r \in R$. In the next iteration all subagents and neighboring strategies sample from these distributions y_{nm}, and the one with the highest expected probability of making a correct suggestion is selected. In case of a referral (i.e. when an agent A_i delegates the choice to another agent A_j), the selection of a subagent a_{jm} is independent of the initial strategy A_i, and the reward is attributed to both A_i and a_{jm}. The dynamics of explorative versus exploitative behavior are realized by adding a cost to such referral activity to make it less or more attractive. The result is a feature vector $V = \langle C, S, Temp, T, D \rangle$ of an ideal suggestion which is selected from the pool of associative images.

The current system-agent is able to make context-dependent decisions, vary its strategy in an explorative and exploitative way, and adapts to the designer's behavior during the interaction process through its online learning capability. However, as highlighted earlier, collaboration is not only dependent on the technical feasibility of making a good decision within a creative space, but also on the interactive behavior strategies. In our next steps, we will therefore focus on the aforementioned dimensions: the explainability and understandability of each agent's choices, whether the system follows or diverges from the designer's choices, and the degree of system assertiveness that benefits a collaborative process.

Conclusion

Using Human-Computer Collaboration to explore, develop and solve underspecified creative problems offers large potentials by extending and augmenting the cognitive, analytical, creative and resource-finding capabilities of all agents involved. In order to exploit these collaborative benefits we need to explore new system approaches that acknowledge the need for evolving objectives as presented in this paper. Finally, more research is needed to identify the relevant interaction behaviors that will enable computing systems to fulfil an equal role in such creative processes.

References

1. Licklider, J. C. Man-computer symbiosis. *IRE Trans. Hum. Factors Electron.* 4–11 (1960).
2. Gray, B. Collaborating: Finding common ground for multiparty problems. (1989).
3. Bozeman, B., Fay, D. & Slade, C. P. Research collaboration in universities and academic entrepreneurship: the-state-of-the-art. *J. Technol. Transf.* **38,** 1–67 (2013).
4. Dell'Era, C. & Verganti, R. Collaborative strategies in design-intensive industries: knowledge diversity and innovation. *Long Range Plann.* **43,** 123–141 (2010).
5. Xue, X., Shen, Q. & Ren, Z. Critical review of collaborative working in construction projects: Business environment and human behaviors. *J. Manag. Eng.* **26,** 196–208 (2010).
6. Allen, J. E., Guinn, C. I. & Horvtz, E. Mixed-initiative interaction. *IEEE Intell. Syst. Their Appl.* **14,** 14–23 (1999).
7. Eckert, C. & Stacey, M. Sources of inspiration: A language of design. *Des. Stud.* **21,** 523–538 (2000).
8. Tekin, C. & van der Schaar, M. Distributed online learning via cooperative contextual bandits. *IEEE Trans. Signal Process.* **63,** 3700–3714 (2015).
9. Agrawal, S. & Goyal, N. Analysis of thompson sampling for the multi-armed bandit problem. in *Conference on Learning Theory* 39–1 (2012).