Conversational Equations: A Conversational Question-Answering Dataset Grounded in Scientific Equations

Anirudh Sundar, William Gay, Shaleen Parikh, Larry Heck

Georgia Institute of Technology

Goal

Create a **Virtual Research Assistant** that

- Amplifies human research
- Is capable of **contextual dialogue**
- Interprets document-grounded **equations**
- Supports **conversational Q-A**

Method

This work introduces a **new dataset**
CONVERSATIONAL EQUATIONS (cEQNS)

- Multi-turn
- Equation-grounded
- Conversational question-answer pairs
- From scientific documents

Preliminary Results

	ROUGE 1	ROUGE 2	ROUGE L
Equation + Question	0.178	0.043	0.115
Equation + Context + Question	0.173	0.043	0.111

Next Steps

- Further ablations
- Dataset augmentation with GPT generated dialogue
- Dataset release

Dataset Collection

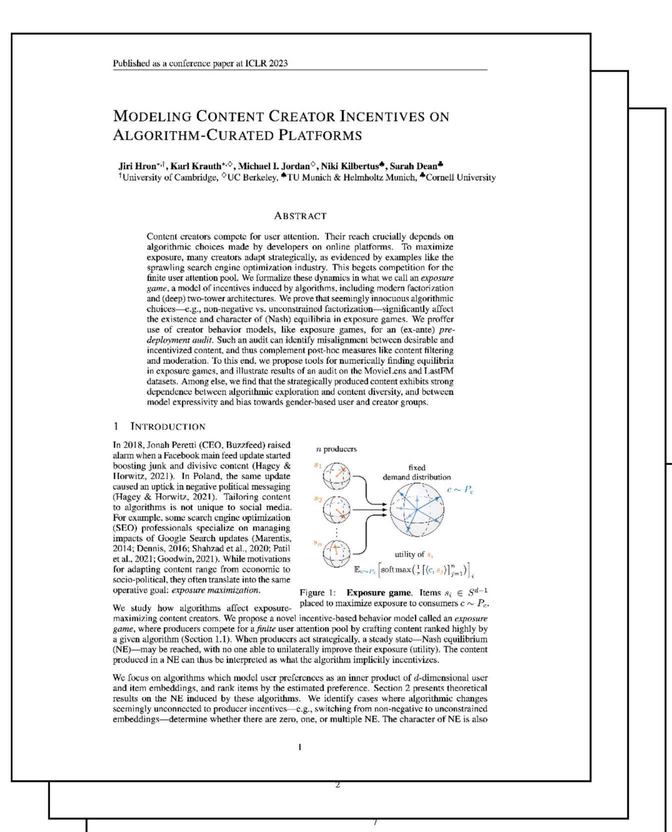

Official Review of Paper6317 by Reviewer 6hAx

Strengths:
1.The paper is well-written and flows very smoothly. Despite the paper's modest space, it provides the required context and necessary explanation in a good job
2.Most parts of the paper have clear motivations. Many of my questions are well answered in the paper.

Weaknesses/Questions:
1.In eq. (2), the paper introduced the temperature parameter τ to control the spread of exposure probabilities. Many of the follow-up analyses and experiments are built on this. However, most recommender systems didn't include this in their objectives. Could you please explain the rationale or connectivity between these?
2.For the ϵ-LNE solver in eq. (4), is this the paper originally proposed or adapted from others? I'm not familiar with the NE-based methods. It seems it's also updating the model with the gradient. Could you please illustrate a little more on the common points and difference between this one and the normal gradient descent method used to optimize ML model?
3.On page 3, I'm confused about the full control assumption. What's the difference between full contry and partial control and why it can abstract away the explicit model of producer actions? Is it possible to list some concrete examples?

Rebuttal

Re motivation for temperature parameter : The temperature parameter can be understood as a relaxation of top-1 recommendation (corresponding precisely to $\tau = 0$). The softmax case ($\tau > 0$) captures nondeterminism, allowing non-zero probabilities of exposure for items with rank larger than 1.

Re solver in experiments: We are not the first to use gradient ascent for finding ϵ-LNE, as the same ideas have been present and used, e.g., in [4, 5] . Essentially, the algorithm runs n *independent* gradient ascent optimizers, each following the gradients of the utility $u_i(s)$. The optimizers execute steps simultaneously, i.e., the iterate at step $t+1$ is obtained by assuming that the other players play the strategy from step t.

Re producer definition in experiments: Defining the exposure game only requires embeddings of the *users* which we fit from rating data using matrix factorization methods (the corresponding item embeddings are not used in the experiments). Independent of the original dataset, we define a number of producers. We focus on *pure* ϵ-LNE, corresponding to each producer creating a single new item.We will clarify this relationship between producers and items in the final revision.

Question 1:
1.In eq. (2), the paper introduced the temperature parameter τ to control the spread of exposure probabilities. Many of the follow-up analyses and experiments are built on this. However, most recommender systems didn't include this in their objectives. Could you please explain the rationale or connectivity between these?

Answer 1:
The temperature parameter can be understood as a relaxation of top-1 recommendation (corresponding precisely to $\tau = 0$). The softmax case ($\tau > 0$) captures nondeterminism, allowing non-zero probabilities of exposure for items with rank larger than 1.

Question 2:
2.For the ϵ-LNE
solver in eq. (4), is this the paper originally proposed or adapted from others? I'm not familiar with the NE-based methods. It seems it's also updating the model with the gradient. Could you please illustrate a little more on the common points and difference between this one and the normal gradient descent method used to optimize ML model?

Answer 2:
We are not the first to use gradient ascent for finding ϵ-LNE, as the same ideas have been present and used, e.g., in [4, 5] . Essentially, the algorithm runs n *independent* gradient ascent optimizers, each following the gradients of the utility $u_i(s)$. The optimizers execute steps simultaneously, i.e., the iterate at step $t+1$ is obtained by assuming that the other players play the strategy from step t.

Equation

$$p_i(c) = \frac{\exp(\tau^{-1}\langle c, s_i\rangle)}{\sum_{i'=1}^{n}\exp(\tau^{-1}\langle c, s_{i'}\rangle)} = \text{softmax}\left(\left[\tau^{-1}\langle c, s_{i'}\rangle\right]_{i'=1}^{n}\right)_i$$

$$\nabla_{\theta_i} u_i(s) = \frac{1}{\tau\|\theta_i\|_2}\left(I - s_i s_i^\top\right)\mathbb{E}[p_i(c)(1-p_i(c))c] = \frac{1}{\tau\|\theta_i\|_2}\left(I - s_i s_i^\top\right)\nabla_{s_i} u_i(s)$$

+

Context

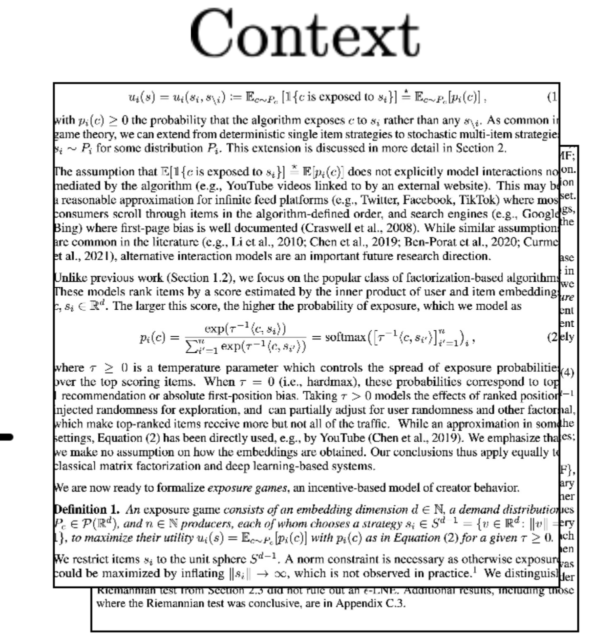

System Architecture

The temperature parameter can be understood as a relaxation of top-1 recommendation (corresponding precisely to $\tau = 0$). The softmax case ($\tau > 0$) captures nondeterminism, allowing non-zero probabilities of exposure for items with rank larger than 1.

Large Language Model

In eq.(2), the paper introduced the temperature parameter τ to control the spread of probabilities. Many of the follow-up analyses and experiments are built on this. However, most recommender systems didn't include this in their objectives. Could you please explain the rationale or connectivity between these?

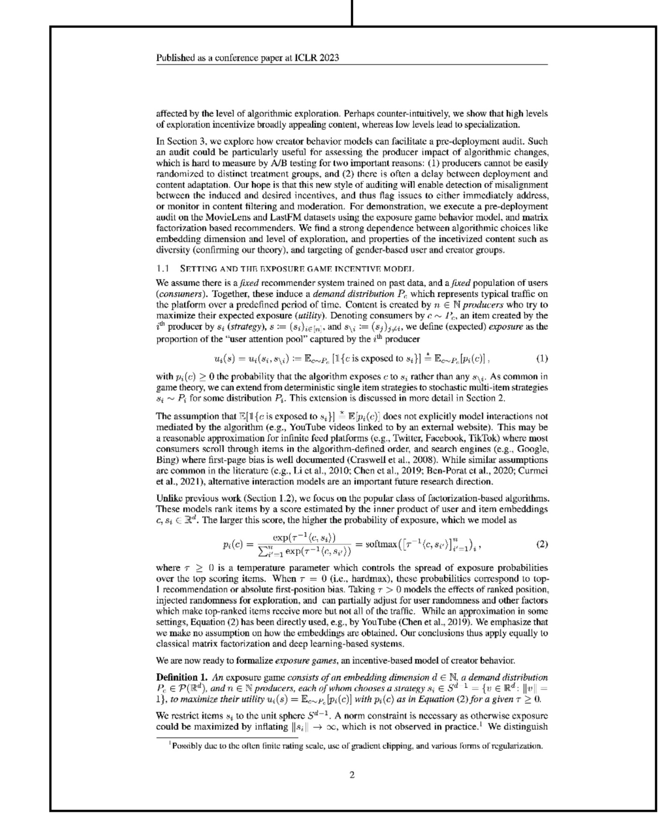

$$p_i(c) = \frac{\exp(\tau^{-1}\langle c, s_i\rangle)}{\sum_{i'=1}^{n}\exp(\tau^{-1}\langle c, s_{i'}\rangle)} = \text{softmax}\left(\left[\tau^{-1}\langle c, s_{i'}\rangle\right]_{i'=1}^{n}\right)_i$$

Scan to be in the loop