Supply-Side Equilibria in Recommender Systems

Jacob Steinhardt

(w/ Meena Jagadeesan and Nikhil Garg)

CLIMB Retreat

Nov 2, 2022

Motivation

Much digital content is mediated by recommender systems:

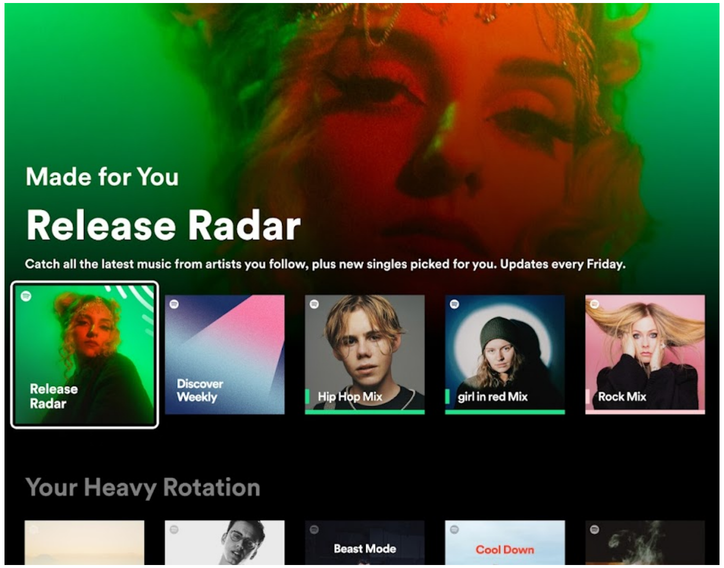

Marketplace where producers **compete** to be recommended to users

Competitive effects: salience, hooks, **genres**

Motivation

Much digital content is mediated by recommender systems:

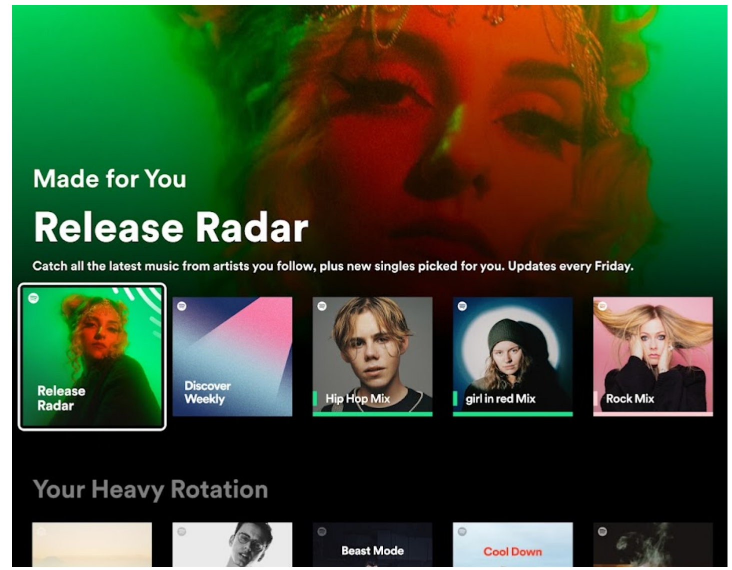

Marketplace where producers **compete** to be recommended to users

Competitive effects: salience, hooks, **genres**

How to understand these?

Formal Setting

Users $u_1, \ldots, u_N$

User vector $u_i \in \mathbb{R}^D_{\geq 0}$ (preferences)

Formal Setting

Users $u_1, \ldots, u_N$

User vector $u_i \in \mathbb{R}^D_{\geq 0}$ (preferences)

P producers. Producer j chooses content $p_j \in \mathbb{R}^D_{\geq 0}$

Formal Setting

Users $u_1, \ldots, u_N$

User vector $u_i \in \mathbb{R}^D_{\geq 0}$ (preferences)

P producers. Producer j chooses content $p_j \in \mathbb{R}^D_{\geq 0}$

Users recommended most similar content p_{j^*}:

$$j^*(u) = \text{argmax}_{j=1}^{P} \langle u, p_j \rangle$$

Formal Setting

Users $u_1, \ldots, u_N$

User vector $u_i \in \mathbb{R}^D_{\geq 0}$ (preferences)

P producers. Producer j chooses content $p_j \in \mathbb{R}^D_{\geq 0}$

Users recommended most similar content p_{j^*}:

$$j^*(u) = \operatorname{argmax}_{j=1}^P \langle u, p_j \rangle$$

Key properties:

- **high-dimensional** decisions
- **heterogeneous** user base

Producer Costs

Producer j wins by making $\langle u, p_j \rangle$ bigger

Trivial solution: take $\|p_j\| \to \infty$

Producer cost function: $c(p_j) = \|p_j\|^\beta$

Profit function:

$$\mathcal{P}(p_j \mid p_{-j}, u_{1:N}) = \underbrace{\left(\sum_{i=1}^{N} \mathbb{I}[j^*(u_i) = j] \right)}_{\text{\# users won}} - \underbrace{c(p_j)}_{\text{cost}}$$

1 Dimension: Compete on Quality

Simplest case: $D = 1$, $N = 1$ user with $u_1 = 1$. $P = 2$ producers

Cost function: $c(p) = |p|^\beta$

1 Dimension: Compete on Quality

Simplest case: $D = 1$, $N = 1$ user with $u_1 = 1$. $P = 2$ producers

Cost function: $c(p) = |p|^{\beta}$

Symmetric mixed equilibrium with c.d.f. equal to $F(p) = p^{\beta/(P-1)}$.

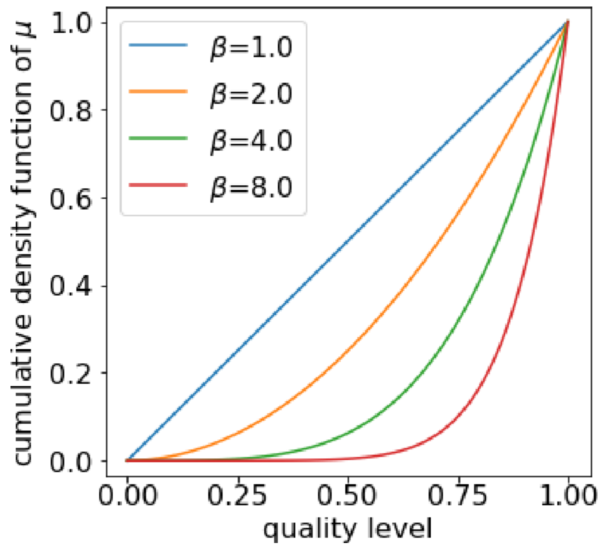

Producers compete on **quality**: more likely to win, but higher cost.

Higher Dimensions: Emergence of Genres

2D example: $u_1 = [1, 0]$, $u_2 = [0, 1]$, $P = 2$ producers.

Density of symmetric mixed equilibrium μ:

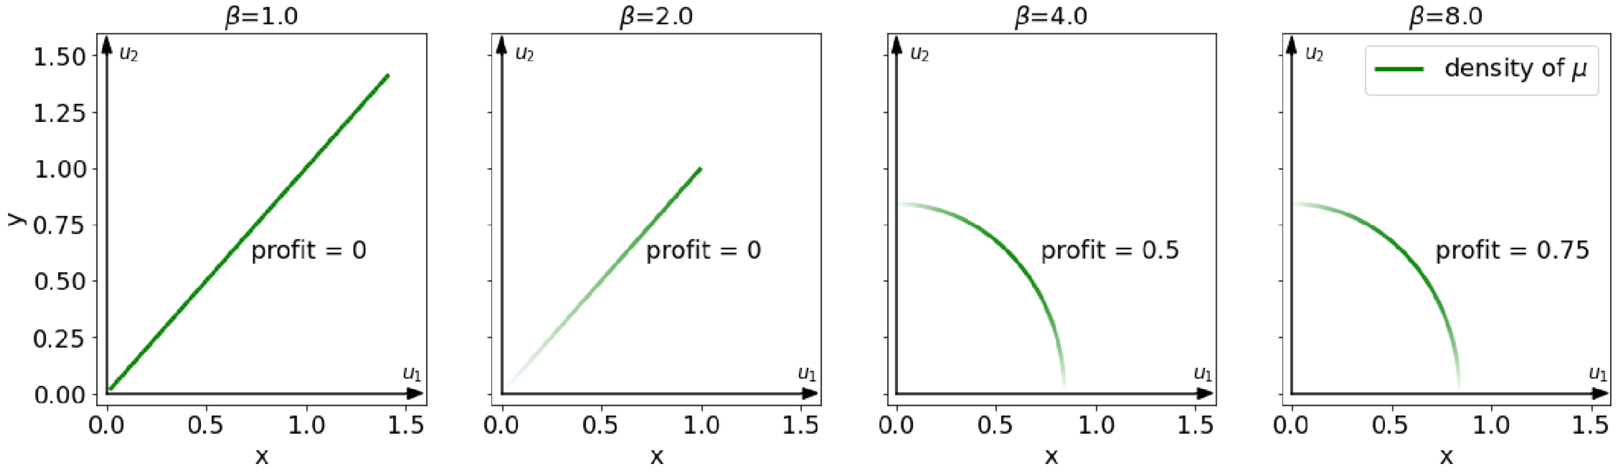

Higher Dimensions: Emergence of Genres

2D example: $u_1 = [1,0]$, $u_2 = [0,1]$, $P = 2$ producers.

Density of symmetric mixed equilibrium μ:

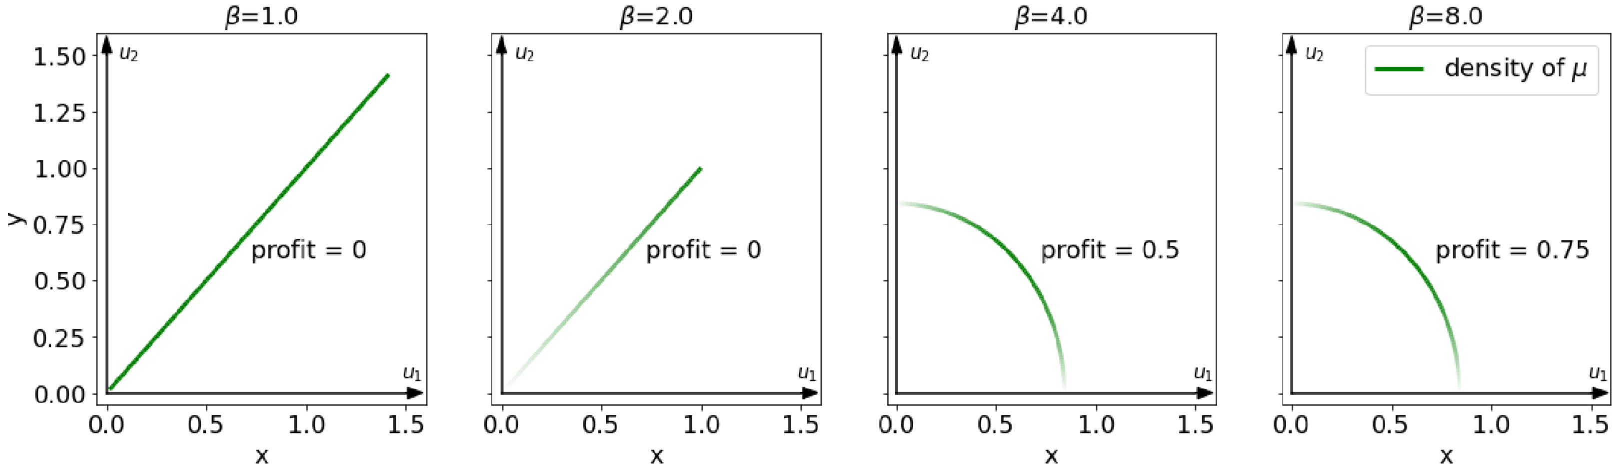

Single genre for $\beta \leq 2$, **infinite genres** for $\beta > 2$.

Producers also achieve **positive profit**.

Higher Dimensions: Emergence of Genres

2D example: $u_1 = [1, 0]$, $u_2 = [0, 1]$, $P = 2$ producers.

Density of symmetric mixed equilibrium μ:

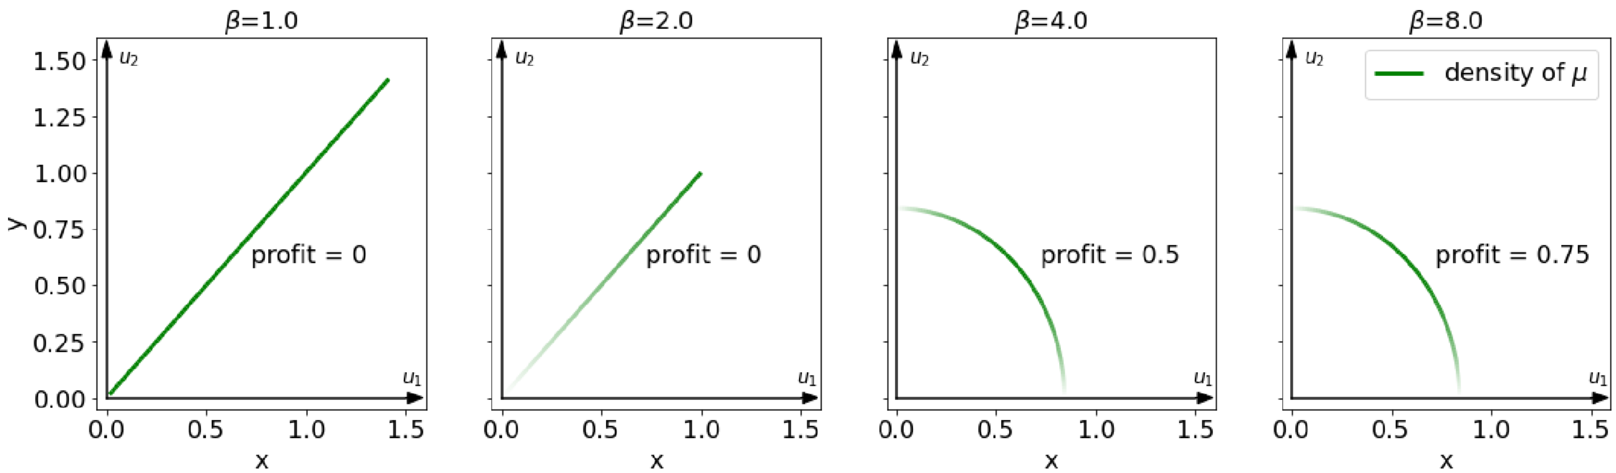

Single genre for $\beta \leq 2$, **infinite genres** for $\beta > 2$.

Producers also achieve **positive profit**.

Will characterize both phenomena!

Genres: Main Result

For a vector p, let $y(p) = (\langle u_1, p\rangle^\beta, \ldots, \langle u_N, p\rangle^\beta)$.

Define $\mathcal{S}^\beta = \{y(p) \mid \|p\| \leq 1\}$.

Genres: Main Result

For a vector p, let $y(p) = (\langle u_1, p\rangle^\beta, \ldots, \langle u_N, p\rangle^\beta)$.

Define $\mathcal{S}^\beta = \{y(p) \mid \|p\| \leq 1\}$.

Theorem (Emergence of genres)

Let $\bar{\mathcal{S}}^\beta$ be the convex hull of $\mathcal{S}^\beta$. Then, all symmetric equilibria are multi-genre iff the following holds: $\max_{y \in \bar{\mathcal{S}}^\beta} \sum_{i=1}^N \log(y_i) > \max_{y \in \mathcal{S}^\beta} \sum_{i=1}^N \log(y_i)$.

In other words, the set $\mathcal{S}^\beta$ must be sufficiently non-convex!

Visualizing S^{β}

$u_1 = [1, 0]$, $u_2 = [0, 1]$:

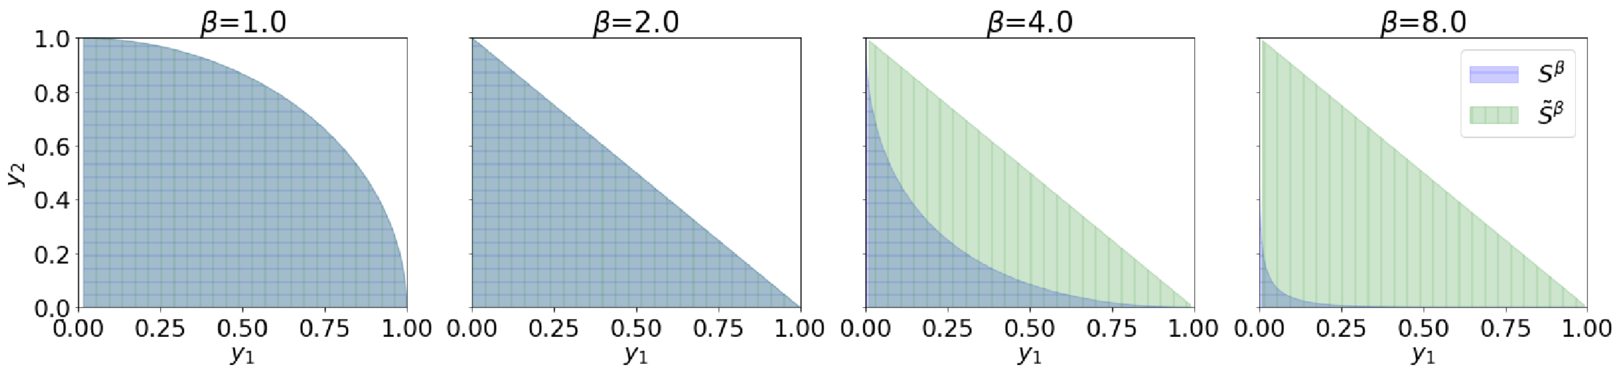

$u_1 = [1, 0]$, $u_2 = [\frac{1}{2}, \frac{\sqrt{3}}{2}]$:

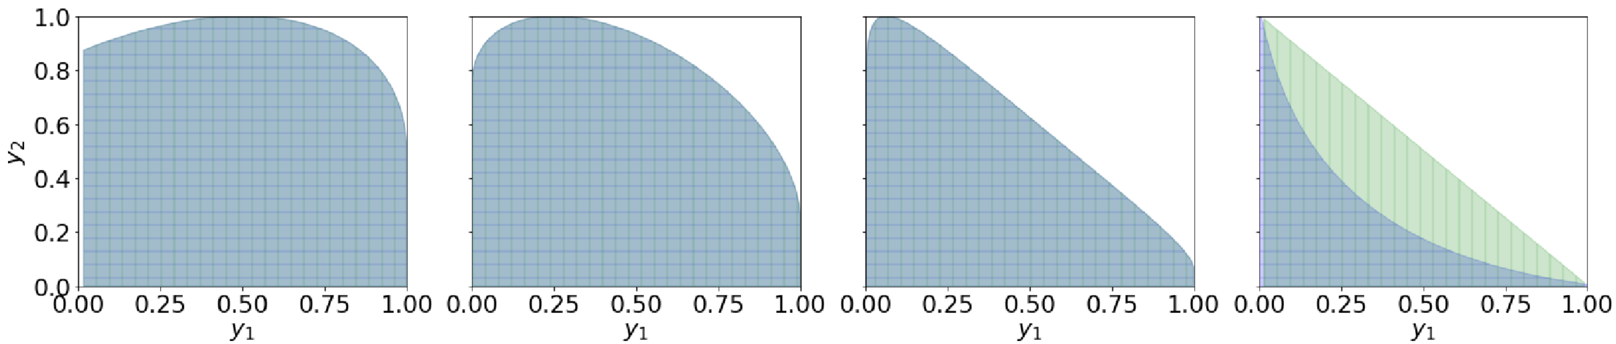

As $\angle(u_1, u_2)$ gets smaller, β threshold gets larger

Corollaries: Structure of Cost

Dependence of genres on user geometry:

Corollary

Let $Z = \| \sum_{i=1}^{N} \frac{u_i}{\|u_i\|_*} \|_*$, where $\| \cdot \|_*$ denotes the dual norm. Then all equilibria are multigenre as long as $Z < N^{1-1/\beta}$.

E.g. random vectors, $Z = O(\sqrt{N})$, multigenre for $\beta = 2 + o(1)$.

Corollaries: Structure of Cost

Dependence of genres on user geometry:

Corollary

Let $Z = \| \sum_{i=1}^{N} \frac{u_i}{\|u_i\|_*} \|_*$, where $\| \cdot \|_*$ denotes the dual norm. Then all equilibria are multigenre as long as $Z < N^{1-1/\beta}$.

E.g. random vectors, $Z = O(\sqrt{N})$, multigenre for $\beta = 2 + o(1)$.

In some cases, $\mathcal{S}^\beta$ is convex regardless of user geometry:

Corollary

(1) If $\beta = 1$, there always exists a single-genre equilibrium μ.
(2) For the ℓ_p-norm and $\beta \leq p$, there always exists a single-genre equilibrium μ.

Producer Profit

Recall 2D example:

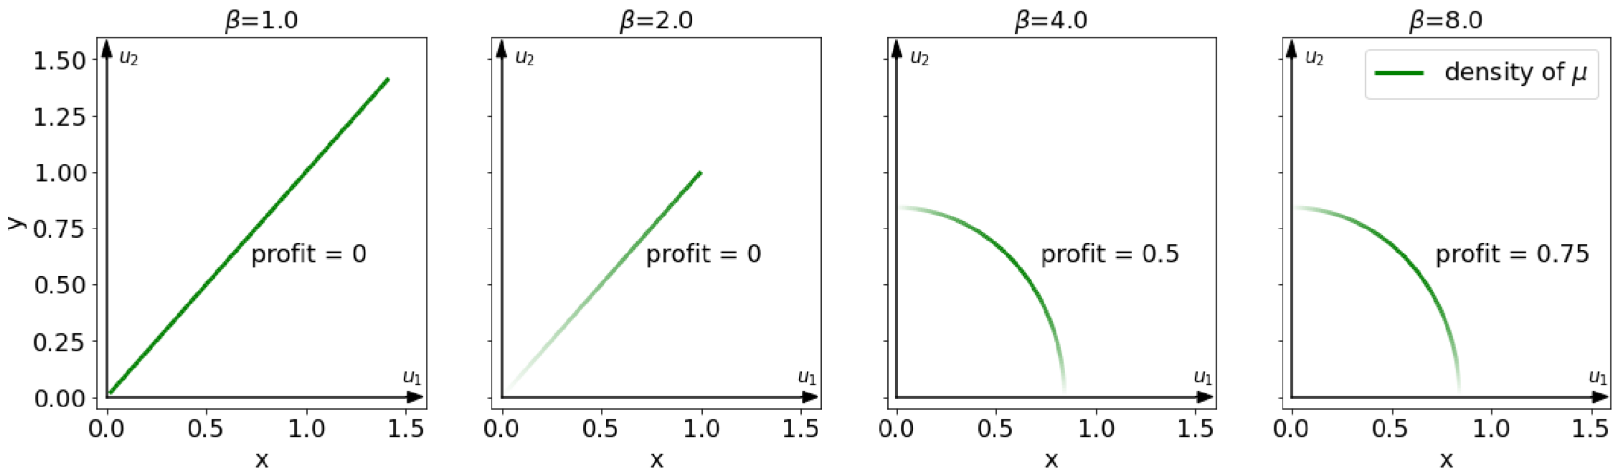

Producers achieve **nonzero profit** (surprising in competitive market)

Seems related to emergence of genres. Can we formalize this?

Producer Profit: Main Result

Theorem (Positive-profit condition)

Let $Q = \max_{\|p\| \leq 1} \min_{i=1}^{N} \langle p, \frac{u_i}{\|u_i\|} \rangle$. Then if $Q < N^{-P/\beta}$, the expected profit at equilibrium is strictly positive.

Interpretation: users more spread out $\rightarrow$ smaller Q $\rightarrow$ easier to profit.

Producer Profit: Main Result

Theorem (Positive-profit condition)

Let $Q = \max_{\|p\| \leq 1} \min_{i=1}^{N} \langle p, \frac{u_i}{\|u_i\|} \rangle$. Then if $Q < N^{-P/\beta}$, the expected profit at equilibrium is strictly positive.

Interpretation: users more spread out $\rightarrow$ smaller Q $\rightarrow$ easier to profit.

Multiple genres necessary for positive profit:

Theorem (1D $\implies$ zero-profit)

Suppose that μ is a single-genre equilibria. Then all producers receive zero expected profit.

Summary

Two key market properties: **high dimensionality** and **heterogeneity**

Two resulting phenomena: **genres** and **positive profit**

Many open questions:

- User utility (vs. number of producers, etc.)
- Noisy recommendations
- Shape of equilibria support

Summary

Two key market properties: **high dimensionality** and **heterogeneity**

Two resulting phenomena: **genres** and **positive profit**

Many open questions:

- User utility (vs. number of producers, etc.)
- Noisy recommendations
- Shape of equilibria support

Thanks!