LyricLens: An Interactive System for Multi-Label Music Content Rating

Kai-Yu Lu
Khoury College of Computer Sciences
Northeastern University
Seattle, WA 98109
lu.kaiy@northeastern.edu

Malhar Sham Ghogare
Khoury College of Computer Sciences
Northeastern University
Seattle, WA 98109
ghogare.m@northeastern.edu

Zihan Su
Khoury College of Computer Sciences
Northeastern University
Seattle, WA 98109
su.zihan1@northeastern.edu

Shanu Sushmita
Khoury College of Computer Sciences
Northeastern University
Seattle, WA 98109
s.sushmita.m@northeastern.edu

Abstract

This paper presents LyricLens, the first online, multi-label lyric classification and rating system. The system's core is a novel classifier, trained on a large-scale Spotify dataset, that identifies four categories of explicit content: Sexual Content, Violence, Language, and Substance Use. LyricLens introduces the Music Content Rating (MCR) framework, a nuanced, five-level certification system (M-E, M-P, M-T, M-R, M-AO) that replaces traditional binary explicit tags. This framework, which also provides specific content descriptors, adapts established rating standards from movies and video games for use with lyrical content. LyricLens is a valuable tool for parental oversight, automated content moderation, and academic research. By providing a detailed, granular analysis of lyrics, the system promotes a more responsible and informed consumption of music, especially for young audiences.

1 Introduction and Background

Music plays a significant role in adolescent development, and popular lyrics often contain Sexual Content, Violence, Substance Use, and strong Language. Greater exposure to degrading sexual lyrics has been linked to earlier sexual activity among adolescents [1], underscoring the need for accurate content assessment and protective measures for young listeners [2]. Evidence from adjacent media also shows that coarse ratings often differ from parental judgments, especially for violent portrayals in film, television, and video games [3]. Within music, early work demonstrates that age ratings are learnable [4], and recent studies move beyond binary labels toward finer granularity [5–7]. While most major streaming platforms rely on a single binary "Explicit" tag, this label is too coarse as it fails to specify the type or intensity of content, fails to capture implicit or metaphorical expressions [8], and ignores frequent co-occurrence, which motivates a multi-label formulation [9, 10].

We address this limitation by introducing LyricLens, the first online, multi-label lyric classification and rating system. Our key contributions are twofold:

- A Multi-Label Classification System: We developed a classifier trained on a large-scale dataset of over 60,000 songs from Spotify. This system is designed to identify four distinct

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: AI for Music: Where Creativity Meets Computation.

and co-occurring categories of explicit content: Sexual Content, Violence, Language, and Substance Use.

- The Music Content Rating (MCR) Framework: Our most significant contribution is the development of the first multi-level music certification system. While movies and video games have established rating systems that include specific content descriptors [11], music has been limited to a binary explicit tag. The MCR framework applies a rule-based system to our classifier's outputs, generating a granular, five-level primary rating (M-E, M-P, M-T, M-R, M-AO) and specific content descriptors. This provides a more detailed and responsible rating, crucial for an audience where young listeners are major consumers.

The interactive web interface for LyricLens supports real-time lyric analysis, providing a valuable tool for parental oversight, automated content moderation on streaming platforms, and academic research into the effects of lyrical content on young audiences. We also release the full dataset, annotation scripts, and the complete code for our classifier and the MCR framework to promote reproducibility and support future research, available at `https://github.com/su1zihan/LyricLens`.

2 LyricLens: System Overview

2.1 Longformer Classifier

The lyric text is preprocessed through a cleaning and lemmatization pipeline before being tokenized by the Longformer model. This fine-tuned model performs multi-label classification to identify four categories: Sexual Content, Violence, Substance Use, and Language. The model was selected for its superior performance, achieving a macro-F1 score of approximately 0.87. The resulting probabilities are then used as inputs for the Music Content Rating framework.

2.2 Music Content Rating

The Music Content Rating (MCR) module converts probability scores from four content categories—Sexual Content, Violence, Language, and Substance Use—into a five-level age-suitability rating, from M-E (Everyone) to M-AO (Adults Only). The system applies an Adults Only rating if a single score exceeds 0.95 or if two or more scores are above 0.85, a condition that captures songs with multiple high-scoring explicit categories. This rule-based mapping is adapted from established media rating standards, including the Motion Picture Association (MPA) film ratings [12] and the Entertainment Software Rating Board (ESRB) system [13]. The module also provides tiered content descriptors for each category, offering granular information on the type and severity of content. Full threshold tables and mapping rules are provided in Appendix B.

2.3 Architecture

LyricLens is an AI-driven system designed to analyze song lyrics for content safety and assign appropriate ratings. The architecture is built for both high performance and user accessibility. A Streamlit-based user interface provides an intuitive platform for lyric input and analysis visualization. The system's robust pre-processing pipeline normalizes and tokenizes text before it is passed to the core component: a fine-tuned Longformer model. This model performs multi-label classification to identify categories like violence and explicit language. A dedicated scoring and rating module then translates the model's probabilities into MCR ratings, ensuring the output is both interpretable and compliant with content standards. The system is engineered with practical features like caching to reduce latency and improve performance, and it uses fallback logic for robust analysis. The entire system is containerized for scalability, making it suitable for large-scale, cloud-based applications. More details provided in Appendix A.

3 Conclusion

LyricLens introduces the first multi-label system for detecting Sexual Content, Violence, Language, and Substance Use in lyrics, and integrates these outputs into the Music Content Rating (MCR) framework. By providing a five-level rating with detailed descriptors, the system supports parental

guidance, enables reliable moderation on streaming platforms, and supplies researchers with an open dataset and codebase for future work.

References

[1] Brian A. Primack, Erika L. Douglas, Michael J. Fine, and Madeline A. Dalton. Exposure to sexual lyrics and sexual experience among urban adolescents. *American Journal of Preventive Medicine*, 36(4):317–323, April 2009. doi: https://doi.org/10.1016/j.amepre.2008.11.011.

[2] Council on Communications and Media. Impact of music, music lyrics, and music videos on children and youth. *Pediatrics*, 124(5):1488–1494, November 2009. doi: https://doi.org/10.1542/peds.2009-2145.

[3] David A Walsh and Douglas A Gentile. A validity test of movie, television, and video-game ratings. *Pediatrics*, 107(6):1302–1308, 2001. doi: 10.1542/peds.107.6.1302.

[4] Anggi Maulidyani and Ruli Manurung. Automatic identification of age-appropriate ratings of song lyrics. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Short Papers)*, pages 583–587, Beijing, China, July 2015. Association for Computational Linguistics.

[5] Linn Bergelid. Classification of explicit music content using lyrics and music metadata. Master's thesis in computer science and engineering, KTH Royal Institute of Technology, Stockholm, Sweden, 2018.

[6] Abraham Albert Bonela, Zhen He, Dan-Anderson Luxford, Benjamin Riordan, and Emmanuel Kuntsche. Development of the lyrics-based deep learning algorithm for identifying alcohol-related words (lydia). *Alcohol and Alcoholism*, 59(2):agad088, mar 2024. doi: 10.1093/alcalc/agad088.

[7] Marco Rospocher and Samaneh Eksir. Assessing fine-grained explicitness of song lyrics. *Information*, 14(3):159, March 2023. doi: 10.3390/info14030159. URL `https://doi.org/10.3390/info14030159`.

[8] Michael Fell, Elena Cabrio, Michele Corazza, and Fabien Gandon. Comparing automated methods to detect explicit content in song lyrics. In Ruslan Mitkov and Galia Angelova, editors, *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 338–344, Varna, Bulgaria, September 2019. INCOMA Ltd. doi: 10.26615/978-954-452-056-4_039. URL `https://aclanthology.org/R19-1039`.

[9] L.K. Dwiyani, I.M.A.D. Suarjaya, and N.K.D. Rusjayanthi. Classification of explicit songs based on lyrics using random forest algorithm. *Journal of Information Systems and Informatics*, 5(2):550–567, 2023. doi: 10.51519/journalisi.v5i2.491.

[10] Hamza Haruna Mohammed, Erdogan Dogdu, Abdül Kadir Görür, and Roya Choupani. Multi-label classification of text documents using deep learning. In *Proceedings of the 2020 IEEE International Conference on Big Data (Big Data)*, pages 4681–4689. IEEE, 2020. doi: 10.1109/BigData50022.2020.9378266.

[11] Cesar A. Perez. A content analysis of the MPAA rating system and its evolution. Honors College Thesis, Georgia Southern University, 2015. URL `https://digitalcommons.georgiasouthern.edu/honors-theses/131`. Available from Georgia Southern Commons.

[12] Motion Picture Association. Classification and rating rules, 2020. URL `https://www.filmratings.com/RatingsGuide`. Effective as revised July 24, 2020.

[13] Entertainment Software Rating Board. Esrb ratings guide. URL `https://www.esrb.org/ratings-guide/`. Accessed August 11, 2025.

A LyricLens: Architecture

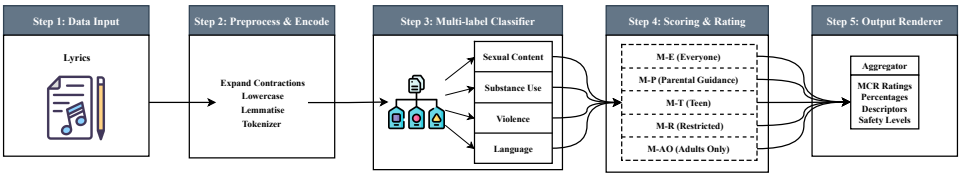

Figure 1: LyricLens Architecture Diagram

A.1 User Interface (Streamlit UI)

This is the front-door to the application, built entirely with the Streamlit framework. It is designed for simplicity and immediate interaction.

- Input: The UI is dominated by a large text area where the user pastes lyrics. This is the primary input mechanism.
- Control: A single "Analyze" button triggers the entire back-end pipeline.
- Display: The interface is split into two columns for a clean presentation. The right-hand column is a dedicated placeholder that dynamically updates to show the analysis results. It uses custom HTML and CSS injected via Streamlit's markdown capabilities to render the polished "Hero Rating Card," content breakdown bars, and safety badges, rather than using standard Streamlit components. This allows for a highly tailored and professional look.

A.2 The Main Pipeline

This is the engine of the application, where the raw text is transformed into a structured rating. It operates as a sequence of specialized processing stages.

A.2.1 Pre-process & Encode:

This is the crucial text cleaning stage. Before the model can analyze the lyrics, the text is rigorously prepared:

- Normalization: Contractions are expanded (e.g., "you're" becomes "you are"), and the text is converted to lowercase.
- Lemmatization: Words are reduced to their root form (e.g., "running," "ran," and "runs" all become "run"). This is more sophisticated than simple stemming and helps the model understand the core meaning without being confused by different word forms.
- Tokenization & Encoding: The cleaned text is broken into tokens (words or sub-words) and then converted into numerical IDs that the Longformer model can understand.

A.2.2 Longformer Inference:

This is the core machine learning step.

- The Model: LyricLens uses a Longformer model, which is a type of Transformer specifically designed to handle long texts. This is critical for song lyrics, which can easily exceed the input limits of other models like BERT.
- The Task: It performs multi-label sequence classification. This means it reads the entire song at once and outputs a set of independent probabilities for four distinct categories: Explicit Language, Violence, Sexual Content, and Substance Use. It does not just classify the song into one bucket; it scores it against all four simultaneously.

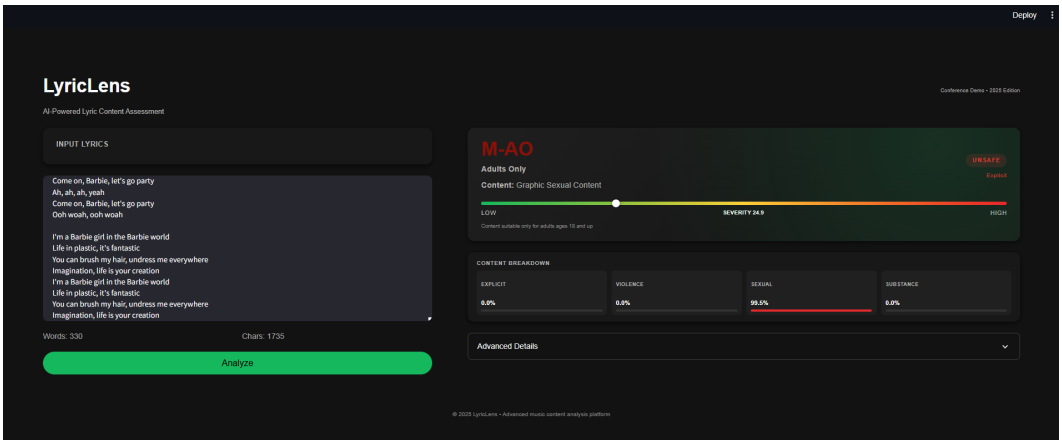

Figure 2: LyricLens Dashboard:Example shows classification of popular song *Barbie Girl* and its appropriate rating. The song has highly sexual content (95%) and therefore is give M-AO rating (Adults Only)

A.2.3 Scoring & Rating:

This stage translates the model's raw, technical output into something meaningful for a human.

- Probability Scaling: The model's output is a set of probabilities from 0.0 to 1.0. These are passed through custom-designed scaling curves that convert them into more intuitive 0-100
- Rating Generation: The final percentage scores are fed into the Music Content Rating (MCR) System. This rule-based module evaluates the combination of scores and assigns the final rating (e.g., M-T for Teen) and selects the appropriate text descriptors (e.g., "Strong Language").
- Content Severity Index: The Content Severity Index synthesizes multi-categorical assessment into a single metric using equal-weight aggregation: $S = \sum(0.25 \times p_i)$ across violence, explicit language, sexual content, and substance use categories. This provides rapid dimensional reduction with balanced assessment, offering intuitive 0-100 interpretability and real-time computational efficiency. The equal weighting prevents category dominance while supporting diverse applications from parental filtering to music analytics research.

A.3 Support Modules

Working in the background, these modules ensure that the application is efficient and reliable.

- Resource Management: This component handles performance optimization. It uses caching to load the large Longformer model into memory only once, making subsequent analyses much faster. It also automatically selects a GPU for processing if one is available, otherwise defaulting to the CPU.
- Safety & Business Logic: This module contains the application rules. It defines the thresholds for safety levels ("Clean," "Moderate" and "Explicit"), manages a fallback mode with simulated data if the main model fails to load, and filters output text to ensure professional presentation.

A.4 Output Renderer

This is the final step in the process. The display_results function acts as the renderer. It takes the aggregated results—ratings, percentages, descriptors, and safety levels—and dynamically constructs the final HTML and CSS to be displayed in the UI's results panel. It is responsible for everything the user sees post-analysis, from the color of the rating pill to the length of the content bars.

B Music Content Rating (MCR) Thresholds and Descriptors

B.1 Primary Ratings

Let $s_{\max} = \max\{s_{\text{sexual}}, s_{\text{violence}}, s_{\text{language}}, s_{\text{substance}}\}$. An M-AO rating is assigned if (i) any $s_i > 0.95$, or (ii) at least two $s_i > 0.85$. Otherwise, the primary rating follows $s_{\max}$, as shown in Table 1.

Table 1: Primary MCR rating thresholds and comparable media standards.

Rating	**Condition on $s_{\max}$**	**Comparable standards**
M-E (Everyone)	all four scores ≤ 0.05	MPA G; ESRB E
M-P (Parental Guidance)	$0.05 < s_{\max} \leq 0.40$	MPA PG; ESRB E10+
M-T (Teen)	$0.40 < s_{\max} \leq 0.70$	MPA PG-13; ESRB T
M-R (Restricted)	$0.70 < s_{\max} \leq 0.95$	MPA R; ESRB M
M-AO (Adults Only)	$s_{\max} > 0.95$ or (≥ 2 scores > 0.85)	MPA NC-17; ESRB AO

B.2 Content Descriptors

Descriptors are assigned per category according to the thresholds in Table 2.

Table 2: Content descriptor thresholds for each category in the MCR system.

Category	**Descriptor**	**Threshold**
Language	Mild Language	$0.05 < s_{\text{language}} \leq 0.40$
	Strong Language	$0.40 < s_{\text{language}} \leq 0.70$
	Explicit Language	$s_{\text{language}} > 0.70$
Violence	Violence	$0.40 < s_{\text{violence}} \leq 0.70$
	Intense Violence	$0.70 < s_{\text{violence}} \leq 0.95$
	Graphic Violence	$s_{\text{violence}} > 0.95$
Sexual Content	Suggestive Themes	$0.40 < s_{\text{sexual}} \leq 0.70$
	Sexual Content	$0.70 < s_{\text{sexual}} \leq 0.95$
	Graphic Sexual Content	$s_{\text{sexual}} > 0.95$
Substance Use	Drug Reference	$0.40 < s_{\text{substance}} \leq 0.70$
	Drug Abuse	$0.70 < s_{\text{substance}} \leq 0.95$
	Glorified Drug Use	$s_{\text{substance}} > 0.95$

C Example Songs for Each MCR Category

We provide one representative song for each MCR category (M-E, M-P, M-T, M-R, M-AO).

C.1 M-E (Everyone)

Song *"Michael"* by The Highwayman

Michael, row the boat ashore, hallelujah
Michael, row the boat ashore, hallelujah
Sister, help to trim the sails (hallelujah
Sister, help to trim the sails (hallelujah)
(Michael, row the boat ashore, hallelujah
Michael, row the boat ashore, hallelujah)
The River Jordan is chilly and cold (hallelujah)
Chills the body but not the soul (hallelujah)
Michael, row the boat ashore, hallelujah
Michael, row the boat ashore, hallelujah
River is deep, and the river is wide (hallelujah)

Milk and honey on the other side (hallelujah)
Michael, row the boat ashore, hallelujah
Michael, row the boat ashore, hallelujah

MCR Rating: M-E: Everyone

C.2 M-P (Parental Guidance)

Song *"Half-breed"* by Cher

My father married a pure Cherokee
My mother's people were ashamed of me
The indians said that I was white by law
The White Man always called me "Indian Squaw"
Half-breed, that's all I ever heard
Half-breed, how I learned to hate the word
Half-breed, she's no good they warned
Both sides were against me since the day I was born
We never settled, went from town to town
When you're not welcome you don't hang around
The other children always laughed at me
"Give her a feather, she's a Cherokee"
Half-breed, that's all I ever heard
Half-breed, how I learned to hate the word
Half-breed, she's no good they warned
Both sides were against me since the day I was born
We weren't accepted and I felt ashamed
Nineteen I left them, tell me who's to blame
My life since then has been from man to man
But I can't run away from what I am
Half-breed, that's all I ever heard
Half-breed, how I learned to hate the word
Half-breed, she's no good they warned
Both sides were against me since the day I was born
Half-breed, that's all I ever heard
Half-breed, how I learned to hate the word
Half-breed, she's no good they warned
Both sides were against me since the day I was born

MCR Rating: M-P: Mild Language

C.3 M-T (Teen)

Song *"We've Got Tonight"* by Bob Seger

I know it's late
I know you're weary
I know your plans
Don't include me
Still, here we are
Both of us lonely
Longing for shelter
From all that we see
Why should we worry?
No one will care, girl
Look at the stars
So far away
We've got tonight
Who needs tomorrow?

We've got tonight, babe
Why don't you stay?

Deep in my soul
I've been so lonely
All of my hopes
Fading away
I've longed for love
Like everyone else does
I know I'll keep searching
Even after today
So there it is, girl
I've said it all now
And here we are, babe
What do you say?

We've got tonight
Who needs tomorrow?
We've got tonight, babe
Why don't you stay?
I know it's late, I know you're weary (weary)
Ooh, I know your plans don't include me
Still, here we are
Both of us lonely
Both of us lonely
We've got tonight
Who needs tomorrow?
Let's make it last
Let's find a way
Turn out the light (turn out the light)
Come take my hand now
We've got tonight, babe
Why don't you stay?

Oh-oh, oh
Why don't you stay?

MCR Rating: M-T: Suggestive Themes

C.4 M-R (Restricted)

"Double Vision" by Foreigner

Feeling down 'n' dirty, feeling kinda mean
I've been from one to another extreme
This time I had a good time, ain't got time to wait
I wanna stick around till I can't see straight
Fill my eyes with that double vision
No disguise for that double vision
Ooh, when it gets through to me, it's always new to me
My double vision gets the best of me

Never do more than I, I really need
My mind is racing, but my body's in the lead
Tonight's the night, I'm gonna push it to the limit
I'll live all of my years in a single minute
Fill my eyes with that double vision

No disguise for that double vision
Ooh, when it gets through to me, it's always new to me
My double vision always seems to get the best of me, the best of me, yeah-ah

Ooh, double vision
I need my double vision
Ooh, (double vision) it takes me out of my head, takin' me out of my head
Ooh, I get my double vision, oh
Ooh, seeing double double, double vision
Ooh, oh my double vision (ooh, double vision)

MCR Rating: M-R: Drug Abuse

C.5 M-AO (Adults Only)

"Barbie Girl" by Aqua

Hiya, Barbie! Hi, Ken! You wanna go for a ride? Sure, Ken! Jump in!
I'm a Barbie girl in a Barbie world Life in plastic, it's fantastic You can brush my hair, undress me everywhere Imagination, life is your creation Come on Barbie, let's go party
I'm a Barbie girl in a Barbie world Life in plastic, it's fantastic You can brush my hair, undress me everywhere Imagination, life is your creation
I'm a blonde bimbo girl in a fantasy world Dress me up, make it tight, I'm your dolly You're my doll, rock and roll, feel the glamour in pink Kiss me here, touch me there, hanky-panky
You can touch, you can play If you say, "I'm always yours" Ooh-whoa-ooh
I'm a Barbie girl in a Barbie world Life in plastic, it's fantastic You can brush my hair, undress me everywhere Imagination, life is your creation
Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh
Make me walk, make me talk, do whatever you please I can act like a star, I can beg on my knees Come jump in, bimbo friend, let us do it again Hit the town, fool around, let's go party
You can touch, you can play If you say, "I'm always yours"
You can touch, you can play If you say, "I'm always yours"
Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh
I'm a Barbie girl in a Barbie world Life in plastic, it's fantastic You can brush my hair, undress me everywhere Imagination, life is your creation
I'm a Barbie girl in a Barbie world Life in plastic, it's fantastic You can brush my hair, undress me everywhere Imagination, life is your creation
Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh Come on Barbie, let's go party Ah-ah-ah, yeah Come on Barbie, let's go party Ooh-whoa-ooh, ooh-whoa-ooh
Oh, I'm having so much fun! Well Barbie, we're just getting started Oh, I love you, Ken!

MCR Rating: M-AO: Graphic Sexual Content