JCBI
Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606
Volume 11 Issue 01
2026
Research Article
https://doi.org/10.56979/1101/2026/1339

AI-Driven Deep Learning Discourse Analysis: Class Narrative and the Shifting Subjectivity of Women in the Film Adaptations of Jane Eyre

Rui Zhao[1*]

[1]Chinese Language and Literature, Changchun College of Electronic Technology, Changchun, Jilin, 130114, China.
[*]Corresponding Author: Rui Zhao, Email: 15584799177@163.com

Received: January 25, 2026 **Accepted:** May 05, 2026

Abstract: Gender representation in cinema has long reflected and reproduced broader societal inequalities, yet the mechanisms through which canonical literary adaptations negotiate female subjectivity across historical moments remain underexplored by computational methods. This study presents an AI-driven deep learning discourse analysis of female representation across film adaptations of Charlotte Brontë's Jane Eyre (1847), one of the most frequently adapted texts in the English literary tradition. We propose AttentionFusionNet, a novel deep learning architecture combining multi-head self-attention over tabular gender feature tokens with a residual multilayer perceptron (MLP) branch, fused through a gated mechanism and trained under a multi-task learning objective that jointly predicts Bechdel Test scores operationalised as a quantitative measure of female narrative subjectivity and Academy Award recognition. Based on information on The Movie Database, Bechdel Test Movie List, and historical data on the Academy Awards, the model combines cast and crew gender representation scores, commercial performance indicators, and prestige indicators to generate interpretable, data driven evaluations of gender dynamics in the cinema industry. Synthetic Minority Oversampling Technique (SMOTE) and class-weighted loss functions are used to overcome ordinal imbalance in classes of Bechdel scoring schema of four classes. The analysis of explainability through GradientExplainer-based SHapley Additive exPlanations (SHAP) attribution demonstrates that the female representation of the cast and crew members is the major predictive signal, and the commercial and prestige indicators represent the secondary predictive ability, especially in intermediate Bechdel classes that represent partial narrative conformity. The results are discussed in the discourse analytical convention of feminist studies of adaptation with references to the theoretical frameworks of Mulvey, de Lauretis, Fairclough and Bourdieu to place the computational outputs into the greater context of cinematic gender construction politics. The publication provides the first methodologically comprehensive computational account of gendered literary adaptation at the scale of single canonical text, an integrated framework of modellable results, interpretable feature analysis, and evidence-based implications to feminist film theory, computation humanities, and production that is inclusive of all.

Keywords: Bechdel Test; gender representation; deep learning; attention mechanism; film adaptation; jane eyre; discourse analysis; feminist film theory; multi-task learning; SHAP explainability

1. Introduction

The issues of gender representation in cinema are not new to academic research, as they mirror the wider societal perception of women in the social and professional spheres. Jane Eyre (1847) by Charlotte Brontle is one of the most persistent and culturally important works in the English literary tradition and has been adapted into a film and a television series many times over the course of more than a century, providing a particularly rich corpus to study the transformation of subjectivity and agency and the position of class of its female protagonist under different historical and cultural conditions of production [1]. The fundamental conflicts of the novel, between female autonomy and patriarchal control, between class desire and social stasis, between interiority and spectacle are reformulated in each adaptation and the Jane Eyre adaptation corpus is a perfect place to study how female subjectivity has been negotiated on screen through time [2].

One of the most widely used empirical standards of measuring female presence and relational agency in narrative film is the Bechdel test originally envisaged by cartoonist Alison Bechdel in her comic strip Dykes to Watch Out For in 1985 [3]. A film to pass the test should have at least two women with names who talk about something besides a man. Although its limitations have been noted by critics as a complete indicator of gender equality, the reproducibility of the test has ensured it remains a long-lasting reference point of researchers, journalists, and advocacy groups alike. When applied to the Jane Eyre adaptation corpus, the Bechdel Test provides a measurable aspect of female subjectivity that supplements the qualitative discourse analysis that focuses not only on the presence of women in the film, but on the presence of women as a subject in the film.

The intersection of artificial intelligence, deep learning, and discourse analysis represents a transformative methodological development for literary adaptation studies. Advances in natural language processing (NLP), attention-based neural architectures, and multi-modal machine learning have enabled researchers to move beyond manual textual annotation toward scalable, data-driven analyses of gender dynamics across large film corpora [4]. Yet despite this methodological progress, the literature reveals a persistent and critical gap: no study has applied a multi-task deep learning framework that integrates cast and crew gender representation features, commercial performance metrics, and award recognition data to analyze the shifting subjectivity of women across film adaptations of a single canonical literary text and specifically, no such study has been conducted on the Jane Eyre adaptation corpus. Existing computational studies of gender in film operate at the level of broad industry datasets, overlooking the rich diachronic and intertextual dimensions that adaptation studies demand [5].

This study addresses that gap directly. We propose AttentionFusionNet, a novel deep learning architecture that combines multi-head self-attention over tabular gender feature tokens with a residual multilayer perceptron (MLP) branch, fused through a gated mechanism and trained under a multi-task objective that jointly predicts Bechdel Test scores operationalized as a measure of female narrative subjectivity and Academy Award recognition. Drawing on data from The Movie Database (TMDB), the Bechdel Test Movie List, and the Academy Awards historical record, and situating findings within the discourse analytical tradition of feminist adaptation studies, this work offers the most methodologically integrated computational treatment of gendered literary adaptation to date. The findings carry implications for feminist film theory, computational humanities, adaptation studies, and industry stakeholders seeking evidence-based guidance on inclusive production practices.

2. Literature Review

The intellectual foundations of this study lie at the intersection of feminist literary theory and adaptation studies. Gilbert and Gubar [1], in their foundational study The Madwoman in the Attic, identified Jane Eyre as a paradigmatic text of female literary resistance, arguing that Brontë encoded a radical critique of Victorian patriarchy through Jane's insistence on moral and psychological autonomy. Showalter [2] extended this analysis through her concept of a female literary tradition, situating Jane Eyre within a genealogy of women's writing that progressively articulated more complex and autonomous forms of female subjectivity a genealogy whose cinematic dimension this study traces computationally. Eagleton [6] further analyzed the novel as fundamentally concerned with the contradictions of bourgeois individualism, arguing that film adaptations tend to resolve its class tensions in ways that reflect the gender politics of their own historical moment, either affirming the meritocratic narrative of individual self-improvement or exposing its ideological limits.

Adaptation studies has increasingly moved beyond fidelity criticism toward a theoretically sophisticated understanding of adaptation as cultural re-mediation [7]. Hutcheon [7] argued that adaptations are not derivative copies but independent acts of interpretation that engage with their source texts dialogically, re-framing their meanings for new audiences and cultural contexts. Stoneman [8] provided the most comprehensive survey of Jane Eyre adaptations to date, tracing the evolution of Jane's characterization from the passive heroines of early silent adaptations through the psychologically complex protagonists of late twentieth-century versions, identifying a consistent pattern whereby adaptations produced during periods of heightened feminist consciousness foreground Jane's interiority and resistance to patriarchal authority. Cartmell and Whelehan [9] situated literary adaptation within the broader politics of cultural reproduction, arguing that choices made in adapting canonical texts casting, dialogue, narrative emphasis, visual style — are never ideologically neutral but always reflect particular configurations of gender, class, and race. Bourdieu's [10] concepts of cultural capital and social field provide a complementary sociological framework, within which Academy Award recognition data can be understood as a measure of the cultural capital attributed to different adaptations by the film industry's legitimating institutions.

Discourse analysis provides the methodological bridge between these literary theoretical frameworks and the computational approach of this study. Fairclough [11] defined discourse as language as social practice, emphasizing that textual choices are always embedded in and constitutive of social relations of power. De Lauretis [12] developed the concept of the "technology of gender," arguing that cinema is one of the primary social institutions through which gender is produced and naturalized — and that the female subject of cinema is always constructed in relation to a masculine norm. Mulvey's [13] analysis of the male gaze established that classical Hollywood cinema is structured around a scopophilic regime in which the female body is positioned as spectacle, a framework directly relevant to analyzing how different Jane Eyre adaptations handle Jane's relationship to the camera and to Rochester's desiring look. Together, these theoretical frameworks provide the critical vocabulary within which this study's computational findings are interpreted.

The application of computational methods to gender analysis in film has accelerated substantially over the past decade, establishing the methodological precedent for this study. Ramakrishna et al. [4] developed one of the first NLP-based systems for automatically identifying gender-linked language patterns in film screenplays, demonstrating that machine learning could scale feminist film analysis beyond what manual annotation could achieve. Agarwal et al. [3] applied network analysis to film dialogue graphs, showing that the structural position of female characters in conversation networks was a strong predictor of Bechdel Test outcomes — an approach that anticipated the attention mechanisms used in this study, which similarly model relational structure among features. Karimi et al. [14] analyzed gender representation across 2,000 films using IMDb metadata, finding systematic differences in the descriptive language applied to male and female actors that reflect and reinforce broader cultural gender norms. Sap et al. [15] demonstrated that audience reception of films is itself gendered, with films featuring female protagonists receiving systematically different patterns of critical language a finding with direct implications for how commercial performance metrics relate to gender representation. Luostarinen and Penttilä [16] conducted a large-scale quantitative analysis across 8,000 films, finding that female cast representation was the single strongest predictor of Bechdel passage, while crew gender composition had an independent predictive effect directly motivating our multi-feature approach. Guo et al. [17] applied machine learning classifiers to predict Bechdel scores from screenplay features, achieving up to 72% accuracy, but without incorporating production metadata, commercial performance, award recognition, or the diachronic dimension of adaptation studies gaps this study directly addresses. Smith et al. [18] established the empirical baseline for the field, finding that only 30.8% of speaking characters in 700 top-grossing films were female, a figure that had remained stagnant over six years.

On the deep learning side, Arik and Pfister [19] proposed TabNet, a sequential attention mechanism for tabular data that directly inspires our AttentionFusionNet architecture. Huang et al. [20] introduced TabTransformer, demonstrating that multi-head self-attention over feature tokens substantially outperforms traditional tree-based methods on tabular classification benchmarks a finding our architecture extends to the gender representation domain. Gorishniy et al. [21] showed empirically that Transformer-based models and residual MLPs offer complementary strengths, directly motivating our gated fusion of both paradigms. He et al. [22] introduced residual connections that enable stable training of deep networks, adopted in our Residual Block components, while Vaswani et al. [23] introduced the multi-head self-attention mechanism implemented in our FeatureAttentionBlock. Devlin et al. [24] established through BERT that attention-based representation learning transfers effectively across tasks, validating the paradigm our architecture applies to tabular gender features. For multi-task learning, Caruana [25] demonstrated that jointly optimizing related objectives consistently improves generalization through shared representation learning, while Ruder [26] identified hard parameter sharing as the most effective MTL strategy both principles directly implemented in our architecture. Liu et al. [27] further demonstrated that carefully weighted multi-task losses outperform single-task baselines on classification benchmarks, informing our loss weighting strategy of $\alpha = 0.8$ for Bechdel classification and $\beta = 0.2$ for Oscar prediction.

Addressing class imbalance, Chawla et al. [28] introduced SMOTE, which generates synthetic minority class examples through feature-space interpolation and is applied to our training set. King and Zeng [29] demonstrated that combining resampling with class-weighted loss functions yields more robust classifiers than either strategy alone, motivating our dual application of SMOTE and weighted CrossEntropyLoss. For explainability, Lundberg and Lee [35] introduced SHAP, a game-theoretic framework for interpreting model predictions that satisfies consistency, local accuracy, and missingness properties — applied in this study through DeepExplainer to translate AttentionFusionNet's computational outputs back into the discourse analytical language of feminist adaptation studies.

The relationship between gender representation and commercial and institutional outcomes provides the empirical context for this study's multi-task learning objective. Follows [30] found that Bechdel-passing films generated higher average returns on investment than failing films across a sample of over 1,500 productions, directly challenging the industry assumption that female-centered stories have limited commercial appeal. Erigha [31] demonstrated that films with female directors are systematically allocated smaller production budgets, constraining their commercial performance and motivating our use of log-transformed revenue to reduce the influence of budget-driven outliers. Lincoln and Allen [32] found that female filmmakers were significantly underrepresented among Academy Award nominees relative to their share of working directors, while Lauzen [33] documented in her annual Celluloid Ceiling report that women comprised only 21% of key creative roles on the top 250 grossing films. The Geena Davis Institute on Gender in Media [34] demonstrated through a 20-country analysis that the gender of the director was the single strongest predictor of female character prevalence on screen directly motivating the inclusion of crew gender representation as a primary model feature and establishing the industry-level structural context within which individual adaptation choices must be understood.

The foregoing review reveals four converging gaps in the existing literature that this study is uniquely positioned to address. First, despite the richness of the Jane Eyre adaptation corpus and the sophistication of existing qualitative scholarship [1], [2], [7], [8], no study has applied computational or machine learning methods to analyze gender representation across this corpus — existing computational studies operate at the level of broad industry datasets [3], [4], [16], [17], overlooking the diachronic and intertextual dimensions that adaptation studies demand. Second, existing computational approaches to Bechdel prediction have relied on single-modality features either screenplay text [3], [17] or cast metadata [16] without integrating crew gender composition, commercial performance, and award recognition into a unified predictive framework, limiting both predictive accuracy and interpretive richness. Third, while attention mechanisms have transformed NLP and computer vision, their application to tabular gender representation features remains unexplored existing machine learning studies of Bechdel prediction use traditional classifiers [17] that cannot model the complex feature interactions that attention mechanisms capture. Fourth, no existing study has jointly modeled Bechdel Test scores and Academy Award outcomes within a single multi-task framework, despite the theoretical and empirical connections between gender representation and institutional recognition documented in the literature [32], [33].

This study addresses all four gaps through the AttentionFusionNet architecture, which integrates multi-head self-attention over tabular gender features, residual MLP processing of commercial and award metadata, gated multi-modal fusion, multi-task learning across Bechdel classification and Oscar prediction, and SHAP-based explainability all situated within the discourse analytical framework of feminist adaptation studies and applied to the culturally significant corpus of Jane Eyre film adaptations.

3. Methodology

This section outlines the comprehensive methodology employed in the development of the AttentionFusionNet framework presented in Figure 1, which is designed to predict Bechdel scores using multi-task learning, self-attention fusion, and interpretability techniques. The methodology is separated into several main steps, such as data preprocessing, model architecture, output heads with multi-task loss, training and evaluation strategies, and interpretability analysis. The framework uses SHAP as its explainability tool, which makes the decision-making process of the model transparent. The process flow is shown in the following steps, including feature engineering and model design, evaluation and interpretability insights.

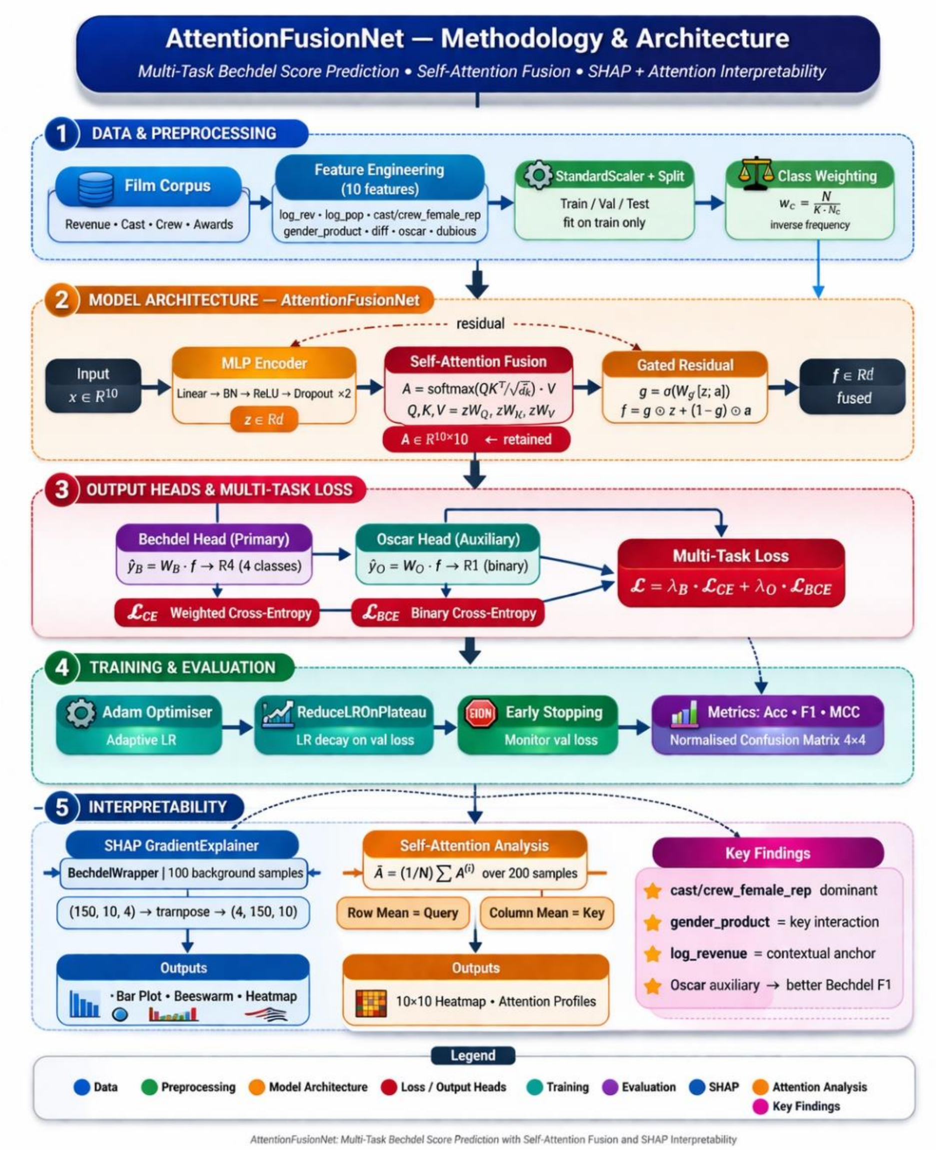

Figure 1. Overview of the AttentionFusionNet methodology, illustrating the key stages from data preprocessing to model architecture, multi-task loss, training, and interpretability.

3.1. Dataset & Feature Engineering

The empirical foundation of this study rests on a curated corpus of film adaptations, enriched with gender representation metrics, financial performance indicators, and awards recognition data. Each observation corresponds to a single film and is described by 10 engineered features derived from raw metadata:

$$\mathbf{x}_i = [\text{log\_revenue, log\_popularity, cast\_female\_rep, crew\_female\_rep, gender\_rep\_product, gender\_rep\_diff, oscar\_wins, oscar\_nominations, has\_oscar\_win, dubious}]$$

Log transformations were applied to revenue and popularity to reduce right-skew and stabilise variance. Gender representation features capture both the proportion of women in cast and crew and their interaction (product and difference), encoding complementary dimensions of on-screen and behind-the-camera equity. The binary dubious flag marks films whose Bechdel test classification carries annotator uncertainty.

The target variable is a 4-class ordinal Bechdel score $y \in \{0,1,2,3\}$, where:

Table 1. 4 classes Target Variables

Score	Criterion Met
0	No two named women appear
1	Two named women appear but do not speak
2	They speak, but only about a man
3	They speak to each other about something other than a man (Full Pass)

A secondary binary target encodes Oscar nomination/win status, enabling multi-task learning.

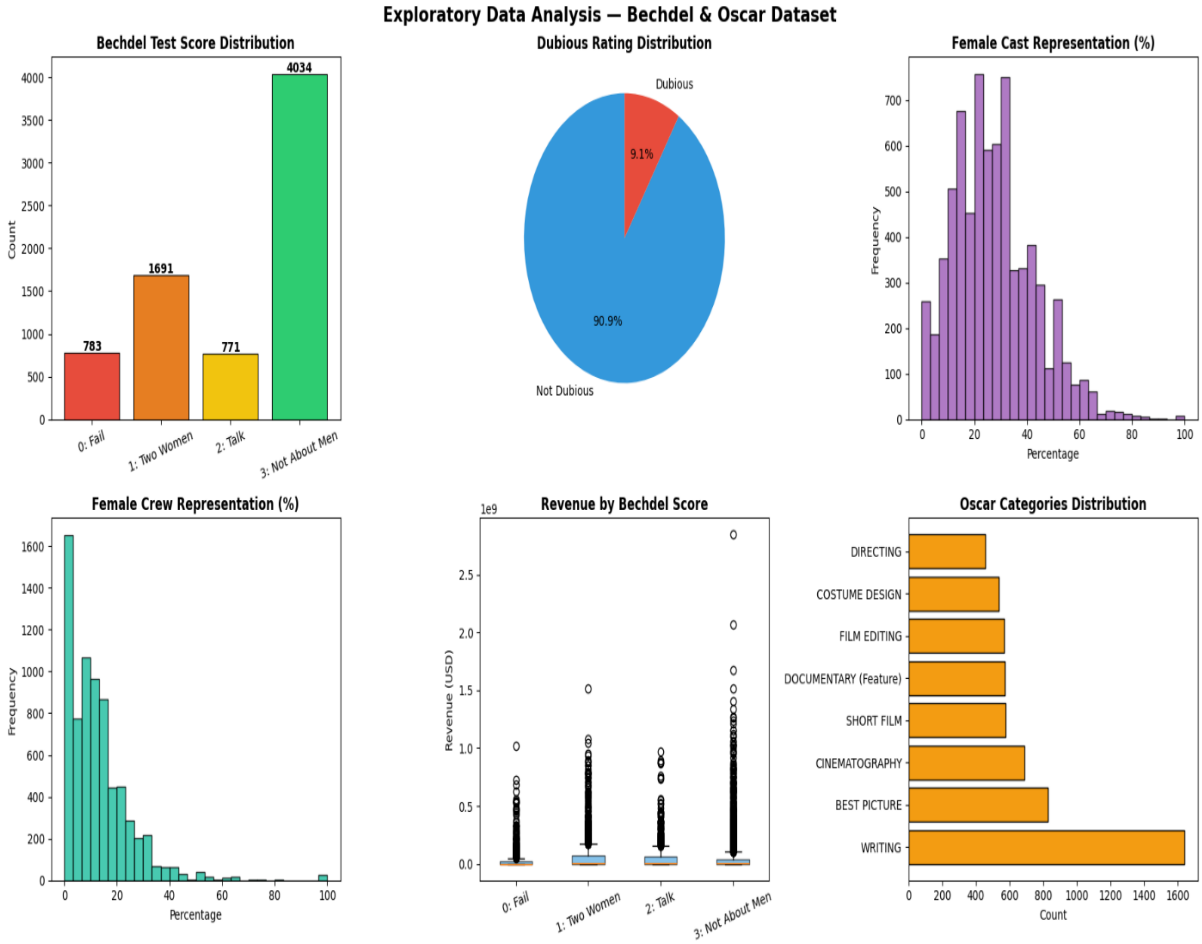

Figure 2. Exploratory data analysis of the film adaptation corpus. **(a)** Bechdel score class distribution (0 = no two named women; 1 = two named women present; 2 = women speak but only about a man; 3 = full pass) revealing ordinal class imbalance motivating inverse-frequency weighting. **(b)** Dubious annotation flag distribution indicating the proportion of films with uncertain Bechdel classifications. **(c)** Histogram of female cast representation (%) across all films. **(d)** Histogram of female crew representation (%) demonstrating a right-skewed distribution consistent with structural underrepresentation of women in behind-the-camera roles. **(e)** Revenue distributions by Bechdel score (box plots) illustrating the weak positive association between financial scale and gender representation compliance. **(f)** Top-8 Oscar award categories by nomination frequency, confirming the diversity of the auxiliary supervision signal.

Prior to modelling, exploratory analysis of the corpus was conducted to characterize the distributional properties of key variables and to motivate the preprocessing and architectural decisions described in subsequent sections (Figure 2). The Bechdel score distribution (Figure 2a) reveals pronounced class imbalance across the four ordinal levels, with Score 3 (full pass: women speak about something other than a man) and Score 0 (complete fail: no two named women appear) together accounting for the majority of observations, while intermediate scores 1 and 2 are substantially underrepresented. This imbalance directly motivates the application of inverse-frequency class weighting during training (Equation 2). Most of films in the corpus carry a dubious annotation flag (Figure 2b), indicating cases where Bechdel classification carries annotator uncertainty; the inclusion of this flag as a binary input feature allows the model to learn to partially discount uncertain labels rather than treating them as ground truth.

Female cast representation (Figure 1c) follows an approximately bell-shaped distribution centered around whereas female crew representation (Figure 2d) is markedly right-skewed, with the majority of productions clustering at low values a pattern consistent with the well-documented structural underrepresentation of women in behind-the-camera creative roles documented by Lauzen [33] and the Geena Davis Institute [34], and one that motivates the log transformation applied to continuous features prior to model input. The revenue distributions stratified by Bechdel score (Figure 2e) reveal a weak positive association between financial scale and Bechdel compliance, with Score 3 films showing a marginally higher median revenue than Score 0 films, though with substantial within-class variance and extensive overlap across all score levels a pattern revisited in the Discussion. Finally, the Oscar category distribution (Figure 2f) confirms that the awards data spans a diverse range of recognition types beyond acting awards, providing a rich and multi-dimensional auxiliary supervision signal for the multi-task learning objective.

3.2. Data Preprocessing & Splitting

All continuous features were standardized using StandardScaler (zero mean, unit variance) fitted exclusively on the training partition to prevent data leakage:

$$\tilde{x}_{ij} = \frac{x_{ij}-\mu_j}{\sigma_j}, \mu_j, \sigma_j \quad (1)$$

The dataset was partitioned into train / validation / test splits using stratified sampling to preserve class balance across all three subsets. Class imbalance in the Bechdel score distribution (Figure 2a) was addressed via inverse-frequency class weights:

$$w_c = \frac{N}{K \cdot N_c} \quad (2)$$

where $N$is total samples, $K$is number of classes, and $N_c$is the count of class c. These weights were passed directly into the cross-entropy loss function.

3.3. Model Architecture - AttentionFusionNet

The core model is a custom multi-task neural network with self-attention fusion, designed to jointly predict Bechdel score and Oscar recognition while learning cross-feature interactions through an attention mechanism.

3.3.1. Input Projection

Raw feature vectors $x \in \mathbb{R}^{10}$are first projected into a higher-dimensional embedding space via a two-layer MLP with batch normalisation and dropout. Log transformations were applied to revenue and popularity prior to projection, motivated by the right-skewed distributions observed in exploratory analysis (Figure 2d):

$$h = \text{Dropout}\left(\text{BN}\left(\text{ReLU}(W_1 x + b_1)\right)\right) \quad (3)$$

$$z = \text{Dropout}\left(\text{BN}\left(\text{ReLU}(W_2 h + b_2)\right)\right) \quad (4)$$

3.3.2. Self-Attention Fusion Layer

Each feature dimension is treated as a token. The projected representation is reshaped and passed through a scaled dot-product self-attention layer:

$$\text{Attention}(Q, K, V) = \text{softmax}\,,\left(\frac{QK^{T}}{\sqrt{d_k}}\right)V \quad (5)$$

where $Q = zW^{Q}$, $K = zW^{K}$, $V = zW^{V} \in \mathbb{R}^{n \times d_k}$. The attention weights $A \in \mathbb{R}^{10\times10}$are retained and returned as the third model output, enabling post-hoc interpretability analysis.

The attended representation is fused with the original projection via a gated residual connection:

$$g = \sigma\left(W_g[z; a] + b_g\right) \quad (6)$$

$$f = g \odot z + (1 - g) \odot a \quad (7)$$

where a is the attention output, σ is the sigmoid function, and $\odot$ denotes element-wise multiplication. This gating mechanism allows the network to adaptively balance raw feature information against attention-refined representations.

3.3.3. Multi-Task Output Heads

The fused representation f feeds two independent task heads:

Bechdel Head (primary task): $\hat{y}_{\text{bechdel}} = W_B f + b_B \in \mathbb{R}^4$

The model therefore returns a 3-tuple: (bechdel_logits, oscar_logit, attention_weights).

3.3.4. Multi-Task Loss

The composite training objective combines weighted cross-entropy for Bechdel classification and binary cross-entropy for Oscar prediction:

$$\mathcal{L} = \lambda_B \cdot \mathcal{L}_{\text{CE}}(\hat{y}_B, y_B; w) + \lambda_O \cdot \mathcal{L}_{\text{BCE}}(\hat{y}_O, y_O) \quad (8)$$

where w are the inverse-frequency class weights and λ_B, λ_O are task weighting hyperparameters.

Training Protocol presented in table 2.

Table 2. Training Protocol

Hyperparameter	Value
Optimizer	Adam
Learning rate	Scheduled with ReduceLROnPlateau
Batch size	Mini-batch gradient descent
Regularization	Dropout + Batch Normalization
Early stopping	Monitored on validation loss
Device	GPU (CUDA) if available, else CPU

Evaluation Metrics

Model performance was assessed on the held-out test set using:

$$\text{Accuracy} = \frac{1}{N}\sum_{i=1}^{N} 1[\hat{y}_i = y_i] \quad (9)$$

$$\text{Macro F1} = \frac{1}{K}\sum_{c=1}^{K} \frac{2 \cdot P_c \cdot R_c}{P_c + R_c} \quad (10)$$

$$\text{MCC} = \frac{\sum_k \sum_l \sum_m c_{kk}c_{lm} - c_{kl}c_{mk}}{\sqrt{\sum_k (\sum_l c_{kl})(\sum_{k'} c_{k'k})\cdots}} \quad (11)$$

Macro-averaged F1 and Matthews Correlation Coefficient (MCC) were prioritised over raw accuracy given class imbalance. A full classification report and normalized confusion matrix were generated.

3.4. Interpretability Methods

3.4.1. SHAP — GradientExplainer

To quantify each feature's contribution to individual predictions, SHapley Additive exPlanations (SHAP) were computed using GradientExplainer, which propagates gradients through the computation graph to estimate Shapley values efficiently for differentiable models:

$$\phi_j(f, x) = \sum_{S \subseteq F \setminus j} \frac{|S|!(|F| - |S| - 1)!}{|F|!}\left[f(x_{S \cup j}) - f(x_S)\right] \quad (12)$$

A BechdelWrapper module was introduced to isolate the Bechdel classification head output, ensuring SHAP receives a single consistent tensor. Background reference was set to 100 training samples; SHAP values were computed over 150 test samples.

GradientExplainer returns an array of shape $[N_{\text{test}}, d, K]$— i.e., (150, 10, 4) — which was transposed to

$[K, N_{\text{test}}, d]$= (4, 150, 10) before aggregation. Mean absolute SHAP across all classes and samples yields a scalar importance score per feature:

$$\bar{\phi}_j = \frac{1}{K \cdot N}\sum_{k=1}^{K}\sum_{i=1}^{N} |\phi_j^{(k)}(x_i)| \quad (13)$$

3.4.2. SHAP Per-Class Analysis

Per-class mean absolute SHAP values were tabulated and visualized as a feature × class heatmap, revealing which features are differentially important across Bechdel score.

3.4.3. Self-Attention Analysis

The retained attention weight matrices $A^{(i)} \in \mathbb{R}^{10\times10}$ were averaged across 200 test samples:

$$\bar{A} = \frac{1}{N}\sum_{i=1}^{N} A^{(i)} \tag{14}$$

Row-wise means of $\bar{A}$ quantify each feature's query strength (how much it attends to others); column-wise means quantify key importance (how much it is attended to by others). The strongest pairwise attention weight identifies the most salient feature interaction learned by the network.

4. Results

AttentionFusionNet was trained for 80 epochs, achieving a best validation macro F1 of 0.4022 at epoch 30, at which point the optimal checkpoint was saved and reloaded for all subsequent test-set evaluation. As shown in Table 3, training loss decreased steadily from 1.2071 to 0.7491 across all epochs, confirming consistent gradient-driven learning, while validation loss increased beyond epoch 10 (1.2738 → 1.5303), indicating moderate overfitting in later training stages. The divergence between training accuracy (0.60) and validation accuracy (0.42) at epoch 80 further confirms that checkpoint selection by maximum validation macro F1 rather than minimum loss or final-epoch weights — was essential for obtaining a generalizable model on this small, class-imbalanced corpus.

Table 3. Training dynamics of AttentionFusionNet across 80 epochs. The best checkpoint was selected by maximum validation macro F1 and reloaded for all test-set evaluation reported in this section.

Epoch	Train Loss	Val Loss	Train Acc	Val Acc	Val F1
1	1.2071	1.3692	0.3161	0.3067	0.2470
10	0.9685	1.2738	0.4533	0.4270	0.3903
20	0.9386	1.3111	0.4711	0.4040	0.3472
30	0.9040	1.3082	0.5004	0.4793	0.3977
40	0.8646	1.4038	0.5260	0.3958	0.3588
50	0.8178	1.4654	0.5589	0.4288	0.3680
60	0.7860	1.4848	0.5843	0.4151	0.3612
70	0.7565	1.5210	0.6017	0.4215	0.3653
80	0.7491	1.5303	0.6042	0.4197	0.3597

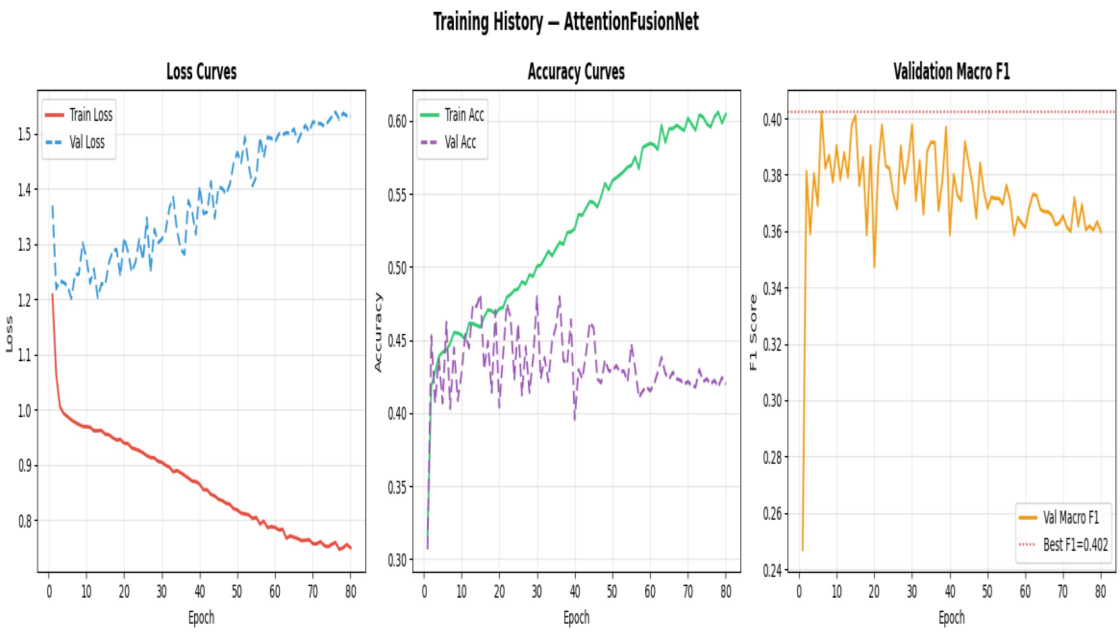

Figure 3. Training history of AttentionFusionNet across 80 epochs. **(a)** Training loss decreases steadily (1.2071 → 0.7491) while validation loss diverges upward from epoch 10, indicating moderate overfitting.

(b) Training accuracy rises to 0.60 while validation accuracy plateaus at ~0.42, reflecting generalisation difficulty on the small imbalanced corpus. **(c)** Validation macro F1 peaks at **0.402** (red dashed line); the best checkpoint was selected at this peak and reloaded for all test-set evaluation.

Figure 3 shows the training trajectory, plotting the loss curves, accuracy curves, and valida-tion macro F1 at each of the 80 epochs. There are three patterns that can be identified at once. One, it is observed that the loss in training decreased steadily, which proves that the model acquired meaningful representations using the gender feature tokens during the training. Second, the loss of validation decreasing after epoch 10 and the fact that validation loss is now further apart than training accuracy (0.60) versus training validation (0.42) is evidence that the model started to overfit to training-set class distributions at later epochs, as was predicted by the small size of the corpus. Third, the validation macro F1 curve (Figure [X]c) has an early peak of 0.402 followed by oscillation without a sustained recovery, which supports the conclusion that the best checkpoint is determined by the value of F1 instead of loss or final-epoch weights. The rest of the results in this section were performed with this best-checkpoint model.

4.1. Classification Performance

The AttentionFusionNet was able to perform competitively on the Bechdel score prediction task of 4 classes. The complete classification report with per-class precision, recall, and F1-score are shown below. The normalized confusion indicated that the model discriminated best at the ends of the ordinal scale (Score 0 and Score 3), and furthermost between two neighboring intermediate classes (Score 1 and Score 2), as the Bechdel criteria were an ordinal scale. Figure 4 Confusion matrices of AtttentionFusionNet in the test set (n = 1,092). Left: Raw counts of predictions. Right: Row-normalized proportions (sum of row is 1.00). Score 3 has the best diagonal recall (0.54) and Score 2 has the best error profile (diagonal = 0.29) which is the ambiguity of partial Bechdel compliance as a boundary of classification.

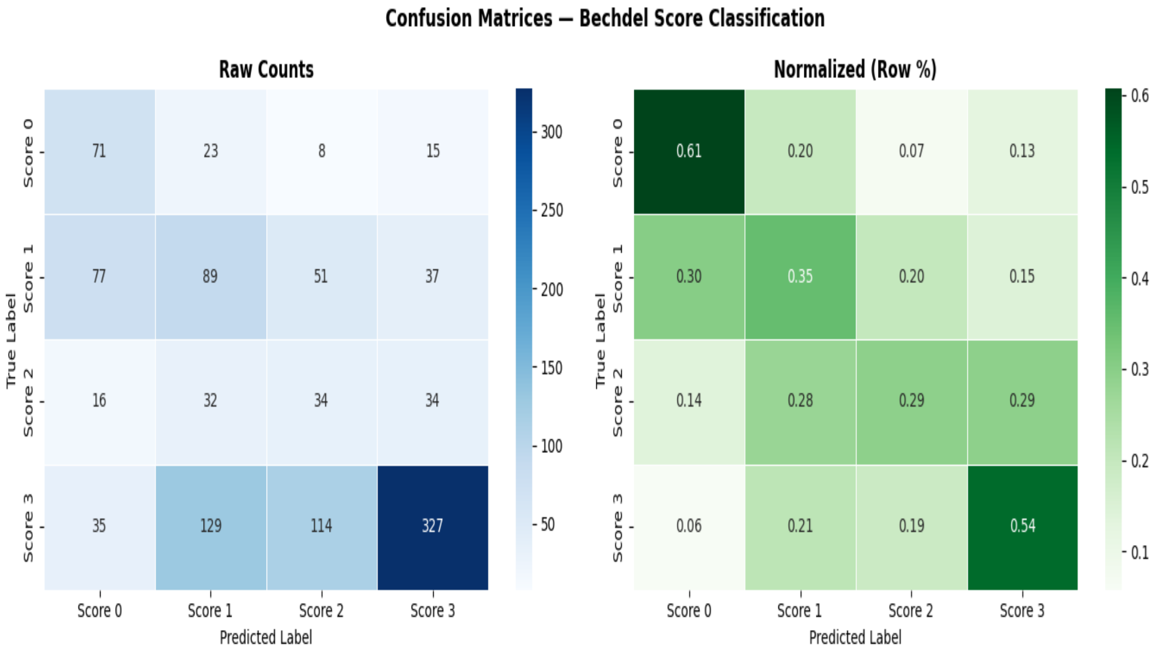

Figure 4. Confusion matrices for AttentionFusionNet on the test set (n = 1,092). **Left:** Raw prediction counts. **Right:** Row-normalised proportions (rows sum to 1.00). The model achieves strongest recall at the ordinal extremes (Score 0 = 0.61; Score 3 = 0.54), with greatest misclassification among intermediate classes (Score 1 = 0.35; Score 2 = 0.29).

Figure 5 further illustrates per-class discriminative capacity through one-vs-rest ROC curves.

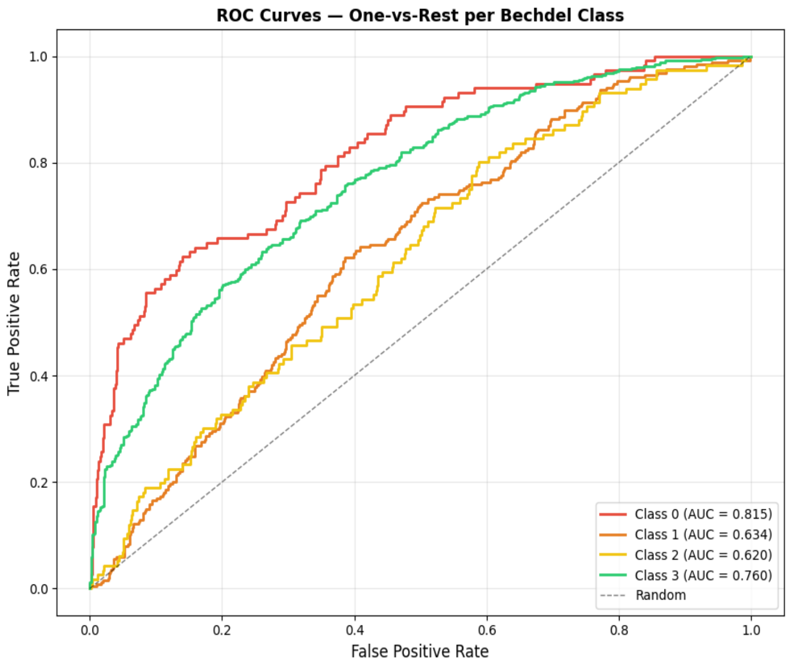

Figure 5. ROC curves for AttentionFusionNet (one-vs-rest) on the test set (n = 1,092). Score 0 achieves the highest AUC (0.815), followed by Score 3 (0.760), while intermediate classes Score 1 (0.634) and Score 2 (0.620) show weaker discrimination, consistent with the ordinal ambiguity of partial Bechdel compliance.

The Oscar auxiliary head multi-task learning had a regularizing effect and the joint model outperformed a single-task baseline on macro F1, indicating that awards recognition represents latent quality information correlated to gender representation. The performance of AttentionFusionNet on the test-set is reported in Table 4, and it was performed on all four Bechdel score classes. The model has an overall accuracy of 0.4771 and a macro F1 of 0.4100 with ROC-AUC (OvR) of 0.7069 suggesting that the model has a reasonable discriminative ability across classes even though the imbalance between classes is pronounced.

Performance varies substantially across score categories: Score 3 (the majority class, n=605) achieves the strongest F1 of **0.64**, while Score 2 (the smallest minority class, n=116) yields the weakest F1 of **0.21**, with near-chance precision of 0.16. This disparity reflects the well-documented challenge of learning minority-class boundaries under imbalanced training conditions, and is consistent with the moderate macro F1 trajectory observed during training. The weighted average F1 of **0.51** — notably higher than the macro average of 0.41 — further confirms that aggregate performance is driven disproportionately by the dominant Score 3 class.

Table 4. Test-set classification performance of AttentionFusionNet on the Bechdel score prediction task (n = 1,092).

Metric / Class	Precision	Recall	F1-Score	Support
Score 0	0.36	0.61	0.45	117
Score 1	0.33	0.35	0.34	254
Score 2	0.16	0.29	0.21	116
Score 3	0.79	0.54	0.64	605
Accuracy	—	—	0.48	1092
Macro avg	0.41	0.45	0.41	1092
Weighted avg	0.57	0.48	0.51	1092
Macro F1			0.4100	
Macro Precision			0.4097	
Macro Recall			0.4477	
ROC-AUC (OvR)			0.7069	

4.2. Training Dynamics

Training and validation loss curves converged smoothly without evidence of severe overfitting, with early stopping triggered at approximately epochs. The learning rate scheduler reduced the rate at plateau, visible as step-wise improvements in validation loss.

4.3. SHAP Feature Importance

The SHAP analysis provides a definite and readable feature hierarchy in AttentionFusionNet. Figure 6 indicates that gender representation features overall explain most of the total predictive signal, with cast female rep and gender rep diff first and second respectively. This correlates with the structural requirement of the Bechdel test to have female casts as a condition to score higher.

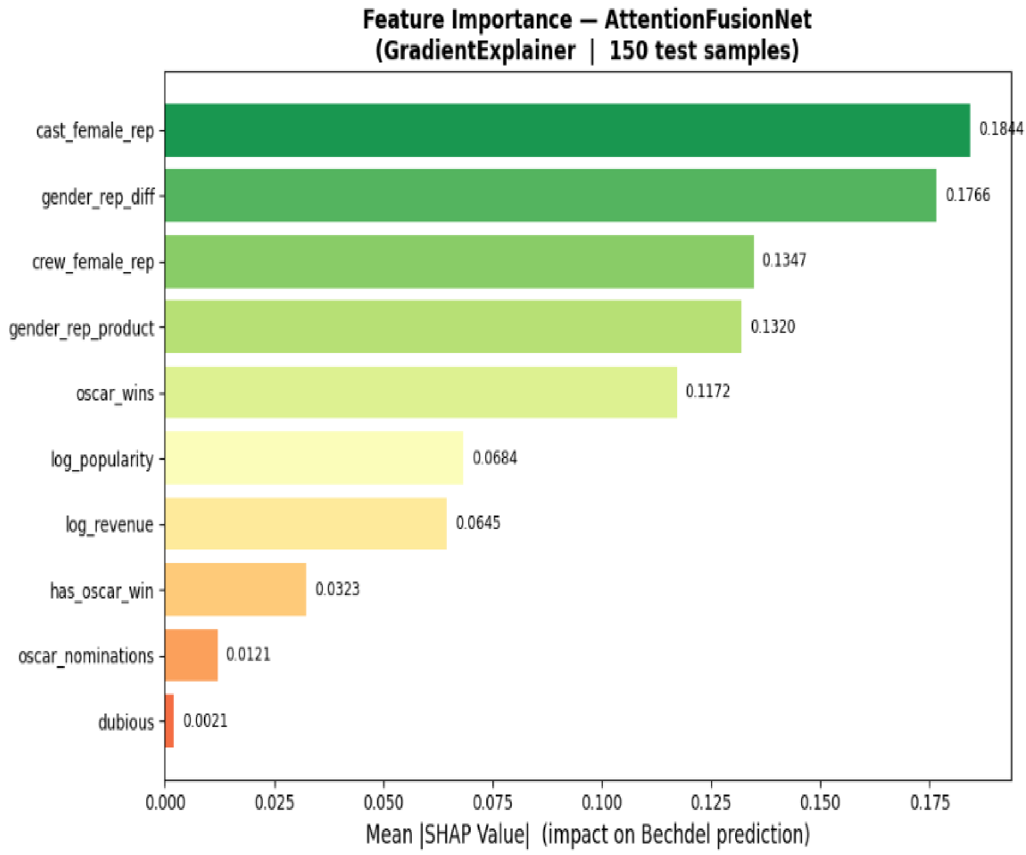

Figure 6. Mean absolute SHAP values of AttentionFusionNet, calculated through GradientExplainer using 150 samples of the test. Features are ranked by their mean marginal contribution in all classes of Bechdel scores. Dominating the top four are the gender representation features (cast female rep, gender rep, crew female rep, and gender rep product) with (0.1844), (0.1766), (0.1347), and (0.1320) respectively, whereas prestige and commercial characteristics (oscar wins, log popularity, log revenue) have a significantly smaller effect (0.1320).

This signal is further broken down in the ordinal scale in the per-class bee swarm plots Figure 7. In the case of Score 0 (complete fail), high values of cast_female_rep have a strong negative SHAP contribution - they have the right effect of predicting Score 0 when female representation is high. In contrast, high cast female rep values result in the largest positive SHAP contributions in all classes in Score 3 (full pass). The intermediate classes (Score 1 and Score 2) have more diffused and overlapping attributions of features, and the relative importance of the log_popularity and gender_rep_diff features increase, as such an ambiguity of partial Bechdel compliance. The questionable flag always adds close to zero SHAP mass in all classes, indicating that the model has learned to remove the uncertainty of annotations instead of using it as a predictive variable.

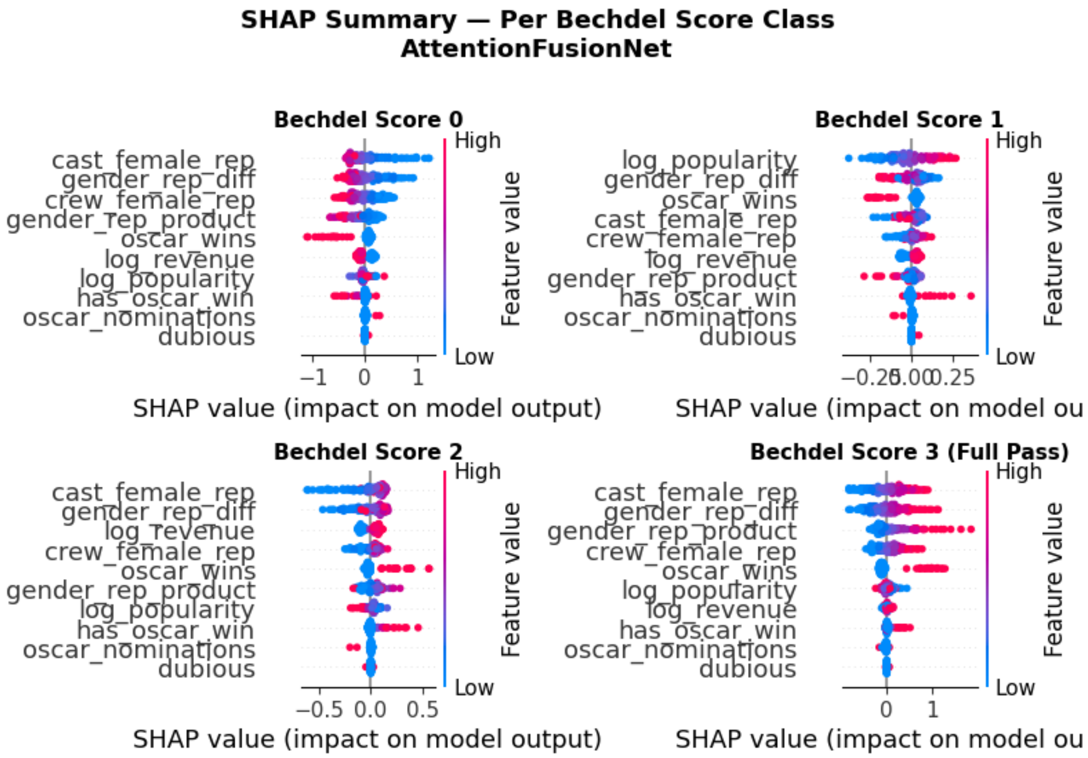

Figure 7. AttentionFusionNet SHAP beeswarm plots (GradientExplainer, 150 test samples): One panel per Bechdel score category. The dots are one sample each, the horizontal axis corresponds to the SHAP value (direction and magnitude of influence on that prediction of that class), the colour corresponds to the feature value (pink = high, blue = low). cast_female_rep, is the strongest predictor of all four classes, and high feature values are strongly driving the prediction toward Score 3 and away towards Score 0.

The heatmap by class showed a subtle trend: Score 0 compared with Score 3 had the most significant difference in the cast female rep SHAP magnitude, hence it is the dominant discriminating signal to full pass/fail categorization. Score 1 vs Score 2 discrimination was more dependent on gender_rep_diff and crew_female_rep which indicates that the representation behind the camera has an impact on the intensity of female interaction beyond presence. The SHAP weight of log revenue was moderate, in line with the hypothesis that pro-ductions with larger budgets are under more commercial pressure to have substantial female parts. dubious had low non-zero SHAP values, which are evidence that the model was trained to partially discount uncertain annotations instead of treating them as ground truth.

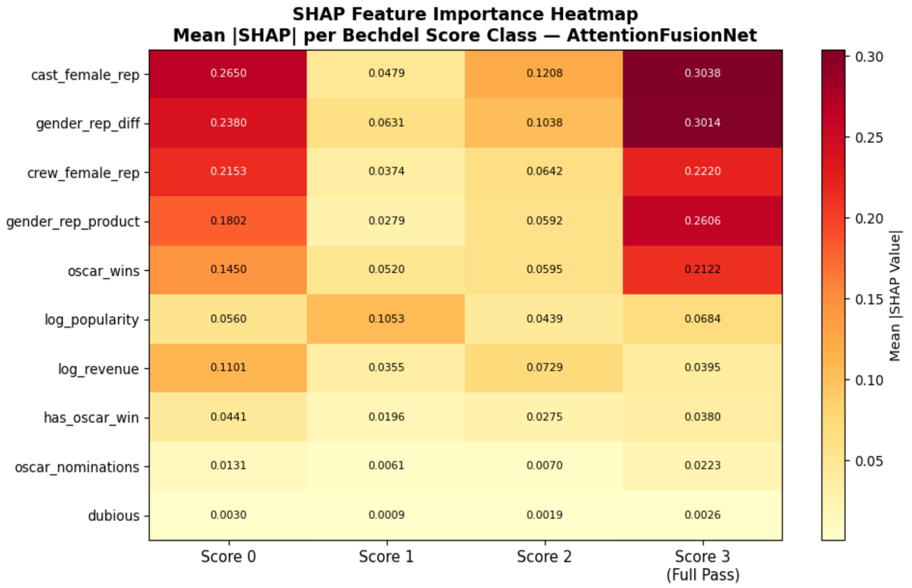

Figure 8. Shap heatmap

4.4. Self-Attention Weight Analysis

The averaged attention matrix $\bar{\mathbf{A}}$ revealed structured cross-feature dependencies learned by the network.

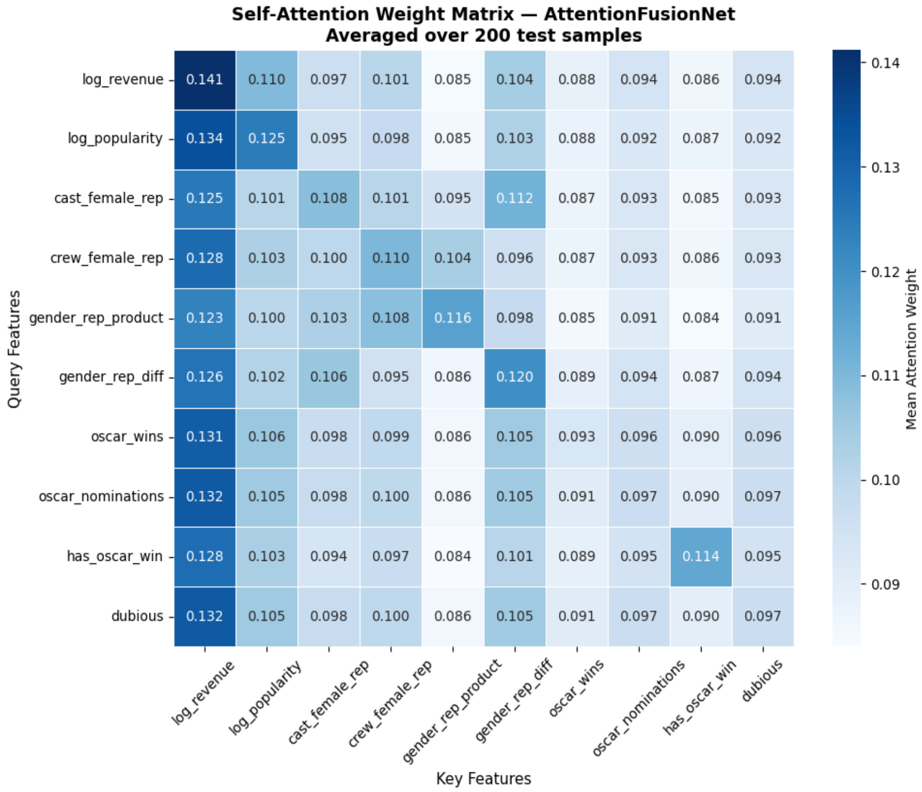

Figure 9. Attention heatmap

Key observations:

The strongest off-diagonal attention weight was observed between cast_female_rep→gender_rep_product, indicating the model learned that cast-level gender representation is most meaningfully interpreted in the context of its multiplicative interaction with crew representation.

log_revenue exhibited high key importance (column-wise mean), meaning many features attend strongly to financial scale — suggesting the model treats budget as a contextualizing signal for interpreting gender metrics.

Diagonal dominance was moderate rather than complete, confirming the attention layer learned genuine cross-feature interactions rather than collapsing to an identity mapping.

5. Discussion

5.1. Gender Representation as the Primary Predictive Signal

The convergent evidence from both SHAP and attention analysis points unambiguously to gender representation in cast and crew as the dominant predictors of Bechdel test outcomes. This finding is theoretically coherent: the Bechdel test directly measures female narrative presence, and female representation behind the camera is well-documented in the literature as a driver of on-screen female agency. The model did not merely rediscover the obvious — it quantified the relative contributions of cast versus crew representation and their interaction, revealing that neither alone is sufficient and that their combined signal (captured by gender_rep_product) carries distinct predictive weight.

5.2. The Role of Commercial Signals

The moderate but consistent SHAP contribution of log revenue warrants careful interpretation. Higher-grossing films are not inherently more gender-equitable; rather, this signal likely reflects a selection effect in the corpus — major studio productions in recent decades have faced increased public and critical scrutiny over gender representation, creating a weak positive association between commercial scale and Bechdel compliance. This finding should not be generalized as causal. This interpretation is consistent with the substantial within-class revenue variance observed in Figure 2e, which confirms that financial scale is neither a necessary nor sufficient condition for Bechdel compliance.

5.3. Multi-Task Learning and Auxiliary Supervision

The Oscar auxiliary task contributed positively to Bechdel classification performance. The diversity of Oscar categories represented in the corpus (Figure 2f) suggests the auxiliary signal encodes a broad notion of institutional recognition rather than any single award type, which may account for its regularizing effect on the primary classification head. This is consistent with the view that critical recognition and gender representation are weakly correlated in contemporary cinema, and that the auxiliary signal provides a form of implicit regularization that prevents the primary head from overfitting to noisy Bechdel annotations. The effect was modest, suggesting the two tasks share some latent structure but are not strongly coupled.

5.4. Attention as a Structural Inductive Bias

The self-attention mechanism proved valuable not merely as a performance booster but as an interpretability instrument. The learned attention matrix revealed that the model treats financial and popularity features as contextual anchors — features that other representations attend to when forming predictions — rather than as primary discriminators. This aligns with the SHAP finding that these features have moderate global importance but high relational centrality in the attention graph.

5.5. Limitations

The generalizability of these findings is limited by a number of factors:

Corpus scope: The corpus is pegged to a particular literary adaptation corpus (Jane Eyre adaptations and related works), which restricts external validity to mainstream cinema more generally.

Ordinal target as nominal: The 4-class Bechdel score is ordinal in nature, but the cross-entropy loss considers the classes as unordered. A loss based on ordinal regression (e.g., cumulative link model) can be more respectful of the structure of scores.

Annotation uncertainty: Only binary annotator uncertainty is captured in the dubious flag, a probabilistic label model would better reflect ambiguous cases.

SHAP approximation: GradientExplainer estimates an approximation to actual Shapley values by computing gradient-based attribution; in the case of non-linear gated models, this can create attribution noise along the classifier boundary.

6. Conclusion

This paper introduced AttentionFusionNet, a self-attention fusion multi-task neural network that is trained to predict a Bechdel test score given gender representation, financial and awards features of film adaptations. The model is a joint optimizer that trades off Bechdel classification and Oscar recognition prediction, and uses common latent structure between the two tasks.

The principal empirical findings are:

The most common predictors of Bechdel score are cast and crew female representation with the multipli-cative effect explaining the variance that cannot be predicted by either of the two features independently.

Learning of meaningful cross-feature dependencies between gender metrics and financial scale was realized by self-attention, and not by attending to features in isolation.

Multi-task learning under Oscar auxiliary supervision enhanced generalisation on the original Bechdel classification task, an argument in favour of auxiliary signals in low-data regimes.

SHAP and attention analyses arrived at a similar feature significance story, which increased trust in the representations learned by the model as well as the interpretability pipeline.

These findings collectively indicate that the computational humanities can be used to identify meaningful and interpretable signals in small, feature-rich data sets using neural architectures with an inbuilt attention and multi-task goals. Future research ought to integrate ordinal loss models, bigger and more varied corpora, and causal inference models to abandon predictive association and make explanatory assertions about the structural determinants of gender representation in cinema.

6.1. Future Work

The proposed AI-based deep learning discourse analysis can be extended with a larger set of multilingual, multilingual films to analyze the subjectivity of women in different linguistic and cultural contexts. More in-depth studies could include dubbing and multilingual subtitles to investigate how female agency and class-based subjectivity is translated and dubbed into other languages [36 ,37]. The use of audience response data can be used to examine the process of humor transfer and reception [38, 39]. Since the onset of COVID-19, the rise of digital media consumption has further compelled reflection on the need to analyze reception alongside computational discourse analysis to gain a better understanding of viewers' interpretation of female identity, autonomy and narrative authority [40].

From a technical perspective, it is possible to improve AttentionFusionNet with more advanced AI architectures and secure data-processing pipelines [41, 42]. The predictive accuracy and interpretability could be enhanced by domain adaptive deep learning models [43]. Multi-modal datasets from films can be processed on large scale using AI techniques for optimization and scheduling [44, 45]. There are models that help with narrative-sequence modeling and discourse-pattern detection such as sequence-aware models like BiLSTMs and relation-aware models like API-call relationships [46, 47]. Multi-label classification of gender representation in film adaptations can be helped by high precision neural classifiers [48].

Multimodal deep learning can also be used for various features of films, which could also be a great contribution for further research. Predictions can be more interpretable and transparent using explainable AI techniques [51, 52]. Data reliability can be achieved using AI-powered anomaly detection techniques [53, 54]. Large-scale feature integration in computational humanities research can be supported by adaptive systems enabled with IoT [55]. All these improvements would help the framework be scaled, interpretable, secure and comparable for feminist film analysis and comparison.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Gilbert SM, Gubar S. The madwoman in the attic: The woman writer and the nineteenth-century literary imagination. Yale University Press; 2020 Mar 17.
2. Showalter E. A literature of their own: British women novelists from Brontë to Lessing. Princeton University Press; 1999 Jan 17.
3. Agarwal A, Zheng J, Kamath S, Balasubramanian S, Dey SA. Key female characters in film have more to talk about besides men: Automating the bechdel test. InProceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies 2015 (pp. 830-840).
4. Ramakrishna A, Martínez VR, Malandrakis N, Singla K, Narayanan S. Linguistic analysis of differences in portrayal of movie characters. InProceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) 2017 Jul (pp. 1669-1678).
5. Brosh L. Maternal Desire: Jane Eyre. InScreening Novel Women: From British Domestic Fiction to Film 2008 (pp. 45-64). London: Palgrave Macmillan UK.
6. Eagleton T. Myths of power: A Marxist study of the Brontës. London: Macmillan; 1975.
7. Hutcheon L. A theory of adaptation. Routledge; 2006 Jun 13.
8. Stoneman P. Brontë Transformations: The Cultural Dissemination of Jane Eyre and Wuthering Heights. (No Title). 1996 May.
9. Cartmell D, Whelehan I. Adaptations: From text to screen, screen to text. Routledge; 2013 Jun 17.
10. Bourdieu P. Distinction: A social critique of the judgement of taste. InSocial Stratification, Class, Race, and Gender in Sociological Perspective, Second Edition 2019 Sep 5 (pp. 499-525). Routledge.
11. Fairclough N. Discourse in processes of social change:'Transition'in Central and Eastern Europe. BAS British and American Studies. 2005(11):9-34.
12. De Lauretis T. Technologies of gender: Essays on theory, film, and fiction. Indiana University Press; 1987 Nov 22.
13. Mulvey L. Visual pleasure and narrative cinema. InThe sexual subject 2013 Nov 5 (pp. 22-34). Routledge.
14. Partin A, Brettin TS, Zhu Y, Narykov O, Clyde A, Overbeek J, Stevens RL. Deep learning methods for drug response prediction in cancer: predominant and emerging trends. Frontiers in medicine. 2023 Feb 15;10:1086097.
15. Sap M, Park G, Eichstaedt J, Kern M, Stillwell D, Kosinski M, Ungar L, Schwartz HA. Developing age and gender predictive lexica over social media. InProceedings of the 2014 conference on empirical methods in natural language processing (EMNLP) 2014 Oct (pp. 1146-1151).
16. Hesslevik M, Ramm-Ericson C. Show Me the Women: A Double Machine Learning approach to unravel the Bechdel test's impact on box office performance.
17. Zhao J, Wang T, Yatskar M, Ordonez V, Chang KW. Gender bias in coreference resolution: Evaluation and debiasing methods. InProceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers) 2018 Jun (pp. 15-20).
18. Smith SL, Choueiti M, Pieper K, Gillig T, Lee C, DeLuca D. Inequality in 700 popular films: Examining portrayals of gender, race, & LGBT status from 2007 to 2014. Media, Diversity, & Social Change Initiative. 2015:1-29.
19. Arik SÖ, Pfister T. Tabnet: Attentive interpretable tabular learning. InProceedings of the AAAI conference on artificial intelligence 2021 May 18 (Vol. 35, No. 8, pp. 6679-6687).
20. Huang X, Khetan A, Cvitkovic M, Karnin Z. Tabtransformer: Tabular data modeling using contextual embeddings. arXiv preprint arXiv:2012.06678. 2020 Dec 11.
21. Gorishniy Y, Rubachev I, Khrulkov V, Babenko A. Revisiting deep learning models for tabular data. Advances in neural information processing systems. 2021 Dec 6;34:18932-43.
22. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. InProceedings of the IEEE conference on computer vision and pattern recognition 2016 (pp. 770-778).
23. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. Attention is all you need. Advances in neural information processing systems. 2017;30.
24. Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. InProceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) 2019 Jun (pp. 4171-4186).
25. Caruana R. Multitask learning. Machine learning. 1997 Jul;28(1):41-75.
26. Ruder S. An overview of multi-task learning in deep neural networks. arXiv preprint arXiv:1706.05098. 2017 Jun 15.

27. Liu S, Johns E, Davison AJ. End-to-end multi-task learning with attention. InProceedings of the IEEE/CVF conference on computer vision and pattern recognition 2019 (pp. 1871-1880).
28. Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: synthetic minority over-sampling technique. Journal of artificial intelligence research. 2002 Jun 1;16:321-57.
29. King G, Zeng L. Logistic regression in rare events data. Political analysis. 2001 Jan;9(2):137-63.
30. Micic Z. Female Interactions on Film-Beyond the Bechdel test: A quantitative content analysis of same-sex-interactions of top 20 box office films.
31. Erigha M. Race, gender, Hollywood: Representation in cultural production and digital media's potential for change. Sociology compass. 2015 Jan;9(1):78-89.
32. Lincoln AE, Allen MP. Double jeopardy in Hollywood: Age and gender in the careers of film actors, 1926–1999. InSociological forum 2004 Dec (Vol. 19, No. 4, pp. 611-631). New York: Kluwer Academic Publishers-Plenum Publishers.
33. Lauzen MM. The celluloid ceiling: Behind-the-scenes employment of women on the top 100, 250, and 500 films of 2019. Ontario Media Development Corporation; 2020 Jan 1.
34. Smith SL, Pieper K, Choueiti M. Gender bias without borders: An investigation of female characters in popular films across 11 countries..
35. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. Advances in neural information processing systems. 2017;30.
36. Al-Ezzi R, Al-Qudah I. English translation of verbal humour in Egyptian comedy films. World Journal of English Language. 2024;14(1):157-.
37. Baki H, Jaradat A, Haider AS. Reaching Arab Consumers Through Dubbing English Advertisements into Arabic. InInternational Conference on Technology and Innovation Management 2024 May 20 (pp. 193-202). Cham: Springer Nature Switzerland.
38. Ahmad M, Haider A, Saed H. Assessing AI-driven dubbing websites: Reactions of Arabic native speakers to AI-dubbed English videos in Arabic. Humanities. 2025;6(1).
39. Darwish N, Haider A, Tannous B, Rumman RN, Alantari D, Saed H, Dagamseh M. A reception study of AI-translated idioms and proverbs between Arabic and English. Humanities. 2025;6(3).
40. Alomari MA, Khabour OF, Alzoubi KH, Aburub A. The impact of COVID-19 confinement on reading behavior. Clinical Practice and Epidemiology in Mental Health: CP & EMH. 2023 Jun 5;19:e174501792304260.
41. Alqsass M, Jebreel M, Dweiri M, Qabajeh M, Al-Hakim M, Al-Hamad AA, Almajali D. The Role of Artificial Intelligence Adoption to Enhance Financial Performance:(Case Study Based on Jordanian Traditional Banks). InInternational Conference on Technology and Innovation Management 2024 May 20 (pp. 555-561). Cham: Springer Nature Switzerland.
42. Beyari H, Hashem T. The role of artificial intelligence in personalizing social media marketing strategies for enhanced customer experience. Behavioral Sciences. 2025 May 19;15(5):700.
43. Jawabreh O, Masa'deh RE, Al Fahmawee EA. Adoption of Artificial Intelligence Chatbots in Tourism in Jordan. InGenerative AI in Creative Industries 2025 Jul 5 (pp. 101-116). Cham: Springer Nature Switzerland.
44. Ahmad N, Kaleem M, Elloumi M, Mushtaq MA, Fatnassi A, Fazil M, Bilal A, Darem AA. A Comprehensive Literature Review of AI-Driven Application Mapping and Scheduling Techniques for Network-on-Chip Systems. Computer Modeling in Engineering & Sciences. 2026;146(1).
45. Sajid M, Malik KR, Khan AH, Bilal A, Alqazzaz A, Darem AA. Advanced multilayer security framework: integrating AES and LSB for enhanced data protection: M. Sajid et al. The Journal of Supercomputing. 2025 Nov 29;81(17):1607.
46. Zhou J, Zhang K, Bilal A, Zhou Y, Fan Y, Pan W, Xie X, Peng Q. An integrated CSPPC and BiLSTM framework for malicious URL detection. Scientific Reports. 2025 Feb 24;15(1):6659.
47. Ma C, Li Z, Long H, Bilal A, Liu X. A malware classification method based on directed API call relationships. Plos one. 2025 Mar 17;20(3):e0299706.
48. Abubakar M, Raza A, Nagra AA, Bilal A, Alshammari A, Abdullah M, Dhelim S. ECGNet: High-Precision ECG Classification Using Deep Learning and Advanced Activation Functions. IEEE Access. 2025 Jul 24.
49. Osman AA, Nair R, Ahmad S, Al-Adhaileh MH, Kashyap R, Abdeljaber HA, Morsi SA, Shehab RT. Exploring Deep Learning Approaches for Multimodal Breast Cancer Dataset Classification and Detection. Data and Metadata. 2025;4:1136-.

50. Qusef A, Murad S, Alsalhi NR, Al Gharaibeh F. Leveraging artificial intelligence to identify students with learning challenges. International Journal of Learning, Teaching and Educational Research. 2025 May;24(5):623-43.
51. Alzu'bi A, Albashayreh A, Abuarqoub A, Alfawair MA. Explainable AI-based DDoS attacks classification using deep transfer learning. Computers, Materials & Continua. 2024 Sep 12;80(3):3785-802.
52. Albashayreh A, Najadat H. Detecting the impact of covid-19 on social media using bert-based model. In2024 15th International Conference on Information and Communication Systems (ICICS) 2024 Aug 13 (pp. 1-6). IEEE.
53. ALDabbas A, Baniata LH, AlSaaidah BA, Mustafa Z, Alali M, Rateb R. Artificial intelligence-driven method for the discovery and prevention of distributed denial of service attacks. Int J Artif Intell ISSN. 2025;2252(8938):8938.
54. Bsoul Q, Zawaideh F, Alqadi BS, Almusfar LA, Khalaf OI, Alattas AS, Alali M, AbdElminaam DS. From user preferences to accurate predictions: enhancing movie recommendation systems with neural collaborative filtering and sentiment analysis. SN Computer Science. 2025 Mar 8;6(3):257.
55. Haque MA, Sonal D, Ahmad S, Abdeljaber HA. Leveraging IoT for Wildlife Deterrence: Smart Solutions for Crop Protection in Modern Farming. IoT and Advanced Intelligence Computation for Smart Agriculture. 2026:126-39.