

BigTokDetect: A Clinically-Informed Vision–Language Modeling Framework for Detecting Pro-Bigorexia Videos on TikTok

Minh Duc Chu¹, Kshitij Pawar¹, Zihao He¹, Roxanna Sharifi²,
Ross Sonnenblick³, Magdalayna Curry⁴, Laura D’Adamo³,
Lindsay Young⁴, Stuart B Murray⁵, Kristina Lerman¹

¹USC Information Sciences Institute

²Keck School of Medicine, USC

³Department of Clinical Psychology, Drexel University

⁴Annenberg School for Communication and Journalism, USC

⁵Department of Psychiatry and Biobehavioral Sciences, UCLA

Abstract

Social media platforms face escalating challenges in detecting harmful content promoting muscle dysmorphic behaviors and cognitions (bigorexia), which disproportionately affects adolescent males and evades moderation by camouflaging as legitimate fitness advice. We address this challenge with BIGTOKDETECT, a clinically informed framework for identifying pro-bigorexia content on TikTok. We introduce BIGTOK, the first expert-annotated multimodal benchmark dataset of over 2,200 TikTok videos labeled by clinical psychiatrists across five categories and eighteen fine-grained sub-categories. Comprehensive evaluation of state-of-the-art vision-language models (VLMs) reveals that while commercial zero-shot models achieve the highest accuracy on broad primary categories, supervised finetuning allows smaller open-source models to achieve better fine-grained subcategory detection. Ablation studies demonstrate that multimodal fusion improves performance by 5-15%, with video features providing the most discriminative signals. These findings support a grounded moderation approach, automating detection of explicit harms while flagging ambiguous content for human review, and establish a scalable framework for harm mitigation in emerging mental health domains.

1 Introduction

Social media platforms face mounting challenges in detecting harmful content that significantly impacts user mental health (Chancellor and Choudhury, 2020; Gorwa et al., 2020). Contemporary platforms like TikTok algorithmically amplify content promoting unrealistic body ideals (Becker, 2004; Minadeo and Pope, 2022), yet much of the harmful material exists in gray areas where legitimate discussions intersect with unsafe messages (Gillespie, 2018). These narratives emerge through complex multimodal combinations, which

include visual imagery, audio cues, and textual descriptions that traditional text-based detection systems do not adequately capture (Kiela et al., 2020a,b). Effective automated moderation requires comprehensive multimodal approaches that can parse nuanced signals across video, audio, and text (Gimeno-Gómez et al., 2024).

Pro-bigorexia content, material promoting muscularity as an aspirational goal, exemplifies these multimodal detection challenges. This content promotes compulsive muscle-building behaviors that disproportionately affect adolescent males (Pope et al., 1997; Mitchison et al., 2022), yet remains severely understudied in computational research compared to traditional eating disorder detection (Chancellor and Choudhury, 2020). Pro-bigorexia material often masquerades as legitimate fitness content through subtle combinations of muscularity displays, extreme workout demonstrations, and coded supplement language (Murray et al., 2017; Kamkari, 2025). The visual similarity between harmful and beneficial fitness content, combined with evolving jargon, renders current detection approaches fundamentally inadequate.

Current automated detection systems face critical limitations while moderating such content. Foundation models trained on general corpora lack domain-specific knowledge to recognize subtle clinical markers and coded language patterns (Murray et al., 2017). Systems relying on keyword filtering and text-based signals miss critical visual and behavioral cues embedded in video content (Kiela et al., 2020b; Gorwa et al., 2020). The dynamic nature of social media, where creators deliberately evolve language to evade detection, poses another challenge to automated approaches (Gillespie, 2018). Effective detection requires integrated analysis of visual behavioral cues, audio transcripts, and textual context, capabilities that remain underexplored in computational mental health research (Gimeno-Gómez et al., 2024).

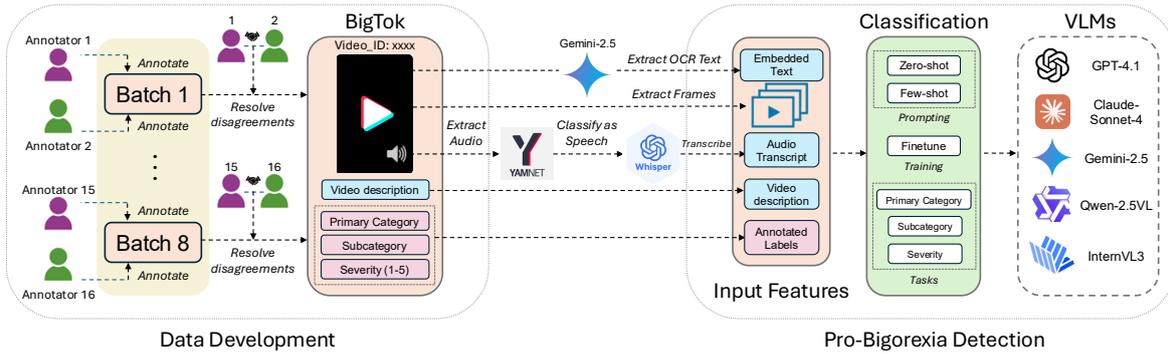


Figure 1: BIGTOK pipeline overview. Left: Expert annotation process with dual annotation and consensus resolution. Right: Multimodal feature extraction (visual, audio, text) and classification evaluation across VLMs using zero-shot, few-shot, and finetuning approaches for primary category, subcategory, and severity prediction tasks.

To address these detection challenges, we make the following contributions:

- We develop the first clinically-informed taxonomy for automated detection of pro-bigorexia content, establishing fine-grained categories spanning body image, nutrition, supplement abuse, exercise practices, and masculinity that enable systematic computational analysis.
- We introduce BIGTOK, a multimodal benchmark dataset of over 2,200 expert-annotated TikTok videos that enables robust evaluation of vision-language models (VLMs) on challenging harmful content detection tasks.
- We construct BIGTOKDETECT, a detection framework that achieves state-of-the-art (SOTA) performance using proprietary and open-source VLMs, with our best models reaching 82.9% accuracy on primary category classification and 69.0% on fine-grained subcategory detection through finetuning.
- We demonstrate through comprehensive ablation studies that multimodal fusion improves detection performance by 5-15% over text-only approaches, with video features providing the most discriminative signals for identifying subtle pro-bigorexia cues.

Figure 1 summarizes our comprehensive approach to multimodal pro-bigorexia detection. Our results demonstrate that clinically-informed VLMs can effectively identify subtle harmful content that evades traditional detection systems. Through systematic evaluation across multiple model architectures and training paradigms, we achieve substantial improvements over existing methods, establishing a new paradigm for detecting nuanced harmful content on social media platforms. The BIGTOKDETECT framework provides both the

computational tools and reproducible methodology needed to address the growing challenge of automated content moderation in mental health domains.

2 Related Work

Muscle Dysmorphic Disorder on Social Media

Muscle dysmorphic disorder, or “bigorexia,” involves preoccupation with insufficient muscularity, driving compulsive behaviors including exercise, rigid dieting, and supplement abuse (Pope et al., 1997; Cooper et al., 2020). Frequent exposure to muscularity-oriented TikTok and Instagram content, including body-transformation videos, supplement promotions, and steroid use, has been linked to elevated rates of probable muscle dysmorphia (Mitchison et al., 2022; Ganson et al., 2025). Hypermasculine online subcultures (e.g., the “manosphere”) intensify these pressures through toxic social comparisons and body-optimization trends (Kamkari, 2025). Despite this prevalence, both psychological research and platform moderation remain focused on anorexia and bulimia, leaving bigorexia largely unmoderated (Lookingbill et al., 2023).

Multimodal Mental Health and Body Image Content Detection

Automated detection of body image content has overwhelmingly targeted anorexia and bulimia (Chancellor et al., 2017; Chancellor and Choudhury, 2020), relying primarily on keyword filters and text-only classifiers despite evidence that vision-language fusion substantially improves detection of subtle harmful imagery (Kiela et al., 2020b). Short-form video platforms have become dominant for youth mental health discourse (Basch et al., 2022), yet eat-

ing disorder detection remains predominantly text-based (Wang et al., 2017; Merhbene et al., 2024). Even dedicated TikTok corpora for eating disorders research exclude muscle dysmorphia labels (Bickham et al., 2025; Donati et al., 2023), while existing qualitative analyses of pro-muscularity forums (Murray et al., 2015) lack the scale needed for robust computational detection. Modern vision-language models—Flamingo (Alayrac et al., 2022), InstructBLIP (Li et al., 2023), GPT-4V (OpenAI, 2023)—enable end-to-end multimodal detection but lack clinically grounded training data for specialized tasks. Existing mental health datasets rely on crowdsourced annotations rather than clinical expertise, potentially missing subtle markers needed to distinguish legitimate fitness content from harmful pro-bigorexia messaging that uses coded language around supplements and steroids (Murray et al., 2017). To our knowledge, no large-scale expert-annotated datasets exist for this increasingly prevalent content type.

3 Pro-bigorexia Content Taxonomy

We develop an empirically grounded taxonomy for pro-bigorexia content characterization, following an iterative four-step process. First, we compile a seed list of keywords from fitness, bodybuilding, exercise, and diet domains, supplemented by terms from foundational muscle dysmorphia research and popular media reports (Cafri et al., 2008; Chu et al., 2024; Willett, 2023). Then, using these keywords, we retrieve 500 TikTok videos over 40 hours, and open-code emergent patterns of bigorexia, such as body-check demonstrations, supplement endorsements, extreme workout narratives, diet tutorials). Next, we draft preliminary categories based on the DSM-5 clinical criteria (American Psychiatric Association, 2013), validated measures such as the Drive for Muscularity Scale (McCreary and Sasse, 2000) and the Muscle Dysmorphic Disorder Inventory (Zeeck et al., 2018). Finally, we convene with clinical experts specializing in eating disorders and body image psychopathologies to refine categories, clarify specialized or coded language (especially around steroids and supplements), and disambiguate edge cases. We iterate to discover new potential categories and keywords.

Our final taxonomy for pro-bigorexia TikTok content consists of five primary categories: *Relationship to Body*, *Relationship to Food*, *Relationship to Exercise*, *Supplements*, and *Masculinity*

(Table 7, Appendix A.1). Each is further divided into clinically grounded subcategories. For each category and subcategory, we provide clear definitions, representative keywords, practical guidelines for recognizing problematic content, and sample videos. The taxonomy facilitates a consistent, multimodal annotation of overt and implicit signals of pro-bigorexia behaviors.

4 BigTok: Expert-Curated Multimodal Dataset Construction

4.1 Data Collection

We source our corpus from TikTok via its official API (TikTok, 2025), querying videos from January 2019 to January 2025 to capture pre- and post-pandemic trends. Guided by domain experts, we curate 40 high-precision query terms mapped to taxonomy subcategories (Table 7, Appendix A.1), retrieving up to 1,000 videos per term, from which we randomly sample 2,400 videos for annotation. To protect user privacy, only video content and captions are exposed to annotators (anonymization steps are described in Appendix C). To ensure robust supervised classification, we construct a negative control group (the *Irrelevant* category) via two mechanisms: (1) proactive sampling using 42 generic trending hashtags (e.g., #viral, #travel) from TikTok’s Creative Center (TikTok Creative Center); and (2) re-assigning videos flagged by experts during annotation as unrelated to bigorexia. This strategy ensures the model learns to distinguish harmful content from both general entertainment and benign daily activities, following established supervised classification practices (Kiela et al., 2020b).

4.2 Annotating Bigorexia Content

Annotator Recruitment and Demographics

We recruit 16 subject matter experts (13 females and 3 males), including licensed clinical psychologists, social workers, and doctoral candidates with research and clinical experience in eating and body image disorders (Table 13, Appendix B.1 for profiles). Their combined expertise in empirical research and direct patient care ensures that our annotations are both theoretically grounded and have clinical relevance.

Annotation Instructions We base our annotation interface on Amazon Mechanical Turk via an invitation-only pool (Appendix B.3). Annotators

are instructed to select the first (mandatory) primary–subcategory pair ($[t_1, s_1]$) and the second (optional) primary–subcategory pair ($[t_2, s_2]$). Every video is also rated for harm on a discrete 5-point Likert scale with 0.5-point increments (1=not harmful and 5=very harmful). The annotation interface is shown in Appendix B.2.

Annotator Training As a pilot study, each annotator independently labels a set of 20–30 videos to validate taxonomy coverage, identify potential edge cases, and familiarize themselves with the annotation interface. We monitor disagreements and encourage detailed note-taking. Annotators then join a group session to discuss flagged issues, collaboratively annotate selected videos while explaining their reasoning, until consensus is reached.

Batchwise Annotation We split the dataset into 8 batches of ~300 videos. For each batch, we filter out videos that are not available (e.g., they have been deleted) or are not in English. Then the batch is assigned to two different annotators, so that every video receives two independent annotations. To assess agreement between a pair of annotators A and B, we compare all combinations of labels and determine the level of agreement:

- **Perfect agreement:** both the first and second primary–subcategory pairs match exactly across annotators, i.e. $\forall i \in \{1, 2\} : [t_i, s_i]^A = [t_i, s_i]^B$
- **Strong agreement:** one primary–subcategory pair matches exactly across annotators, i.e. $\exists i \in \{1, 2\} : [t_i, s_i]^A = [t_i, s_i]^B$
- **Weak agreement:** at least one primary category matches across any labels, but the subcategories differ, i.e. $\exists i \in \{1, 2\} : t_i^A = t_i^B \wedge s_i^A \neq s_i^B$
- **Disagreement:** no common primary category across labels, i.e. $\forall i \in \{1, 2\}, t_i^A \neq t_i^B \wedge s_i^A \neq s_i^B$.

We then flag examples with weak agreement or disagreement on pro-bigorexia detection and examples where severity scores differ by more than two points. The two annotators are invited to a structured discussion to articulate their reasoning, review cases that need consultation, address disagreements, and eventually aim to reach consensus, although consensus is not mandatory. After the discussion, each annotator independently re-annotates the flagged videos. Further details on this procedure are provided in Appendix B.4.

Label Aggregation While each video may receive multiple label pairs (primary-subcategory), to ensure consistent and reproducible evaluation, we assign a single label pair to each video. For videos with perfect agreement, we randomly choose one of the label pairs. For strong agreement, we adopt the shared primary–subcategory pair; for weak agreement, we choose the mandatory (first) label over the optional (second) one; remaining disagreements are resolved by random tie-breaking. The final harm severity score is the average of the two annotators’ ratings. This simplification reflects current limitations of LLMs in reliable multi-label classification (Ma et al., 2025).

We assess inter-rater reliability with Cohen’s κ (Cohen, 1960) for two annotators. Initial agreement is moderate (strong $\kappa = 0.43$ – 0.69 ; weak $\kappa = 0.59$ – 0.83 ; Table 8). After structured consensus discussions and re-annotation, reliability increases substantially (strong $\kappa = 0.58$ – 0.81 ; weak $\kappa = 0.78$ – 0.94), ensuring that our iterative process yields high-quality, consistent labels. All Cohen’s κ values for each batch are shown in Table 8 (Appendix B.5). Severity ratings showed moderate agreement (ICC = 0.413, 69.4% within one point), typical for subjective clinical assessments and highlighting the complexity of harm evaluation in this domain (Figure 6, Appendix B.5).

4.3 Data Statistics

Our BIGTOK dataset originates from a raw core set of 2,210 expert-annotated TikTok videos (median duration 20.6 seconds, 59.8% under 30 seconds). To construct robust classification benchmarks, we apply a strict preprocessing pipeline detailed in Section B.4 (Appendix), which includes removing corrupted or non-English files and supplementing the dataset with verified irrelevant videos sampled via popular keywords (Table 7). Table 1 presents the final train/test split statistics resulting from this processing. Specifically, the primary category benchmark (Task 1) comprises 2,554 videos split into 1,966 training and 588 balanced test samples, while the subcategory benchmark (Task 2) utilizes a relevance-filtered subset of 1,934 videos (1,472 training, 462 test). Within the processed primary categories (excluding *Irrelevant*), *Relationship to Body* (587 videos) and *Relationship to Exercise* (548 videos) represent the largest classes. At the subcategory level, *Muscularity Self-objectification* (204 videos) predominates, reflecting the prevalence of physique-display con-

Primary Category	Train	Test	Total
Relationship to Body	489	98	587
Relationship to Food	310	98	408
Relationship to Exercise	450	98	548
Supplement Abuse	201	98	299
Relationship to Masculinity	66	98	164
Irrelevant	450	98	548
Total	1,966	588	2,554

Subcategory	Train	Test	Total
Muscularity Self-objectification	170	34	204
Leanness Self-objectification	59	34	93
Muscle Dissatisfaction	72	34	106
Rigid Food Rules	144	34	178
Unsafe Food	59	34	93
Cheat Meals	60	34	94
Excessive Exercise	129	34	163
Predebtting Exercise	28	4	32
Maladaptive Coping	82	34	116
Exercise-Induced Functional Impairment	43	6	49
Toxic Motivation	102	34	136
Anabolic Steroids	61	34	95
Legal APEDs	56	10	66
Hormone Therapy	77	34	111
Relationship to Masculinity	130	34	164
Irrelevant	200	34	234
Total	1,472	462	1,934

Table 1: BIGTOK Classification Benchmark statistics showing the final train/test split distribution after data processing.

tent on fitness-oriented TikTok. Examples from the final dataset appear in Tables 14 and 15 in Appendix A.1, with additional annotation statistics provided in Appendix B.5. The raw annotation counts are shown in Table 9 in the Appendix.

5 Classifying Multimodal Pro-Bigorexia Content via VLMs

To develop an automated detection system for pro-bigorexia content, we apply SOTA VLMs to this challenging multimodal classification task. We leverage both proprietary and open-source VLMs through zero-shot prompting, few-shot learning, and finetuning approaches to achieve high performance in identifying subtle pro-bigorexia behaviors across video, audio, and text modalities.

5.1 Task Definition

We define three evaluation tasks for TikTok video classification. **Task 1: Primary Category Classification** involves predicting one of the five primary

categories of pro-bigorexia content (*Body, Food, Supplements, Exercise, Masculinity*), or *Irrelevant*.

Task 2: Subcategory Classification involves predicting the specific subcategory within the selected primary category. **Task 3: Severity Estimation** involves predicting a continuous severity score on a 1-5 scale, where 1 indicates no harm and 5 indicates extreme harm.

For each task, we split the data into training and test sets at a 3:1 ratio using stratified sampling. The test set is strictly balanced across all classes, while the training set is adjusted by downsampling the majority categories to mitigate class imbalance. The sampling procedures and statistics after sampling are provided in Appendix C.

5.2 Models

Given the multimodal nature of our classification task, involving text from captions, audio transcripts, images, and video content, we require SOTA models designed for multimodal data, particularly video understanding. We evaluate three commercial API-based VLMs: GPT-4.1 (OpenAI, 2025) and Claude-Sonnet-4 (Anthropic PBC, 2025) (both process images), while Gemini-2.5-Flash (Comanici et al., 2025) natively accepts video input. We also evaluate two open-source models: Qwen2.5-VL (Qwen Team, 2025), and InternVL3 (Zhu et al., 2025), which uses Qwen2.5 pre-trained base models and Variable Visual Position, which achieves strong performance on video benchmarks.

Model sizes and versions are listed in Table 10, Appendix D.1. Due to computational constraints, we focus on small and medium-sized models. For open-source models, we set $temperature = 0.1$ to encourage deterministic outputs. For commercial API models, we utilize the provider defaults ($temperature = 1.0$). To ensure the validity of this choice, we conducted a robustness check comparing $T = 1.0$ vs $T = 0.1$ for all commercial models (see Appendix E.1). We observe negligible performance differences, confirming that our results are robust to hyperparameter variations. Implementation details and hardware specifications are provided in Appendix D.1. Details about the train and test datasets are in Appendix C.

5.3 Input Features

To capture the full range of pro-bigorexia signals, we extract four complementary input features: **Visual**: We adapt visual inputs to model architectures. For Gemini-2.5-Flash and Qwen2.5-VL, which

support native video processing, we provide the raw video stream. For InternVL3, GPT-4.1, and Claude-Sonnet-4, we treat video as a sequential image task. We extract 4 equally spaced frames per video (unless otherwise noted in ablation studies) to balance temporal coverage with context-window constraints. **Audio:** We classify audio using YAM-Net (Hershey et al., 2017) and transcribe detected speech via Whisper (Radford et al., 2022). To prevent bias, we retain all videos regardless of audio type: 40.2% contain intelligible speech, while the remaining 59.8% (music or ambient noise) are assigned empty transcripts. **On-Screen Text:** We use Gemini-2.5-Flash to extract overlaid text, creators’ captions and annotations, since many TikTok creators rely on on-screen text for key messages. **Caption:** We include the original TikTok description (user caption and hashtags) from the video metadata.

5.4 Training and Inference Paradigms

Zero-Shot Prompting We create prompts for zero-shot classification by incorporating taxonomy definitions of the primary and subcategories and instructing the model to select the appropriate category as the label. The full zero-shot prompt templates for Tasks 1 and 2 are shown in Figure 8 and 9, Appendix D.2. We evaluate the models on balanced test sets of 588 examples for primary classification (Task 1) and 466 examples for subcategory classification (Task 2), with each class evenly represented (Appendix C)

Few-Shot Prompting For few-shot experiments, we sample a fixed set of in-context examples from the training data: two videos per primary category (12 examples total) for Task 1 and one video per subcategory (16 examples) for Task 2 (prompt details are in Figures 10 and 11, Appendix). Each example is presented with its full multimodal features: video frames, audio transcript, on-screen text, caption, and the corresponding label. We evaluate the VLMs on a balanced test set of 462 examples (Appendix C).

Finetuning Due to cost constraints, we limit finetuning to open-source VLMs. We instruction-tune various model variants (up to 72B) on our annotated datasets. The batch size is 8, and the learning rate is $5e-5$. The inference hyperparameters are similar to zero-shot and few-shot prompting (temperature = 0.1 and top_p = 0.9).

5.5 Pro-Bigorexia Classification Results

Table 2 shows that while large commercial models like Claude-Sonnet-4 achieve the highest zero-shot performance on broad primary categories, parameter-efficient finetuning allows smaller open-source models (e.g., InternVL3-8B) to surpass them on fine-grained subcategory detection by capturing domain-specific clinical nuances.

Primary vs. Subcategory Detection Commercial API-based models continue to lead in broad-category detection: Claude-Sonnet-4 (zero-shot) attains the highest accuracy of 0.829 and F1 of 0.827. Although few-shot open-source models (e.g., InternVL3-38B) reach competitive levels (0.806 accuracy), none yet surpass the commercial zero-shot benchmark on primary categories. However, the trend reverses for fine-grained subcategory detection tasks. InternVL3-8B (finetuned) achieves the top subcategory F1 of 0.675, outperforming commercial methods, which plateau around 0.640 F1. This demonstrates that while large-scale pre-training confers an advantage for general classification, parameter adaptation better captures nuanced subtype distinctions.

Impact of Model Scale Our experiments reveal a divergence in finetuning effectiveness relative to model scale. While finetuning dramatically boosts smaller models (7B–8B), for instance, raising Qwen2.5-VL-7B subcategory accuracy from 0.383 (zero-shot) to 0.684 (finetuned), it frequently degrades the primary classification performance of larger open-source models (32B–38B). For example, InternVL3-38B drops from 0.784 (zero-shot) to 0.767 (finetuned) on primary categories. This likely stems from overparameterization under data scarcity: our ~1,500 training samples create a capacity mismatch for larger models, leading to overfitting, whereas 7B–8B models achieve a more optimal bias-variance tradeoff.

Impact of Prompting Strategies Few-shot prompting yields modest boosts in primary-category performance for some models (e.g., GPT-4.1 improves from 0.796 to 0.813 accuracy), highlighting the value of in-context learning for broad tasks. However, in subcategory detection, few-shot gains are minimal or inconsistent. This suggests that for fine-grained tasks, a small set of exemplars may act as noise, anchoring the model to idiosyncratic patterns rather than clarifying the complex boundaries between subtypes. By contrast, super-

Model	Training	Primary Category				Subcategory			
		Acc.	P _m	R _m	F1 _m	Acc.	P _m	R _m	F1 _m
GPT-4.1	Zero-shot	0.796	0.808	0.792	0.792	0.652	0.639	0.503	0.532
	Few-shot	0.813	0.820	0.813	0.813	0.680	0.679	0.643	0.639
Claude-Sonnet-4	Zero-shot	0.829	0.832	0.829	0.827	0.670	0.679	0.636	0.640
	Few-shot	0.819	0.829	0.818	0.818	0.665	0.674	0.627	0.632
Gemini-2.5-Flash	Zero-shot	0.805	0.807	0.805	0.805	0.666	0.665	0.616	0.621
	Few-shot	0.776	0.782	0.776	0.775	0.663	0.645	0.613	0.621
Qwen2.5-VL-7B	Zero-shot	0.539	0.744	0.539	0.526	0.383	0.554	0.327	0.306
	Few-shot	0.689	0.733	0.689	0.689	0.236	0.149	0.142	0.140
	Finetuning	0.784	0.810	0.784	0.776	0.684	0.687	0.669	0.667
Qwen2.5-VL-32B	Zero-shot	0.733	0.788	0.733	0.742	0.556	0.637	0.495	0.513
	Few-shot	0.804	0.812	0.804	0.807	0.654	0.662	0.594	0.607
	Finetuning	0.776	0.800	0.776	0.775	0.658	0.662	0.647	0.645
InternVL3-8B	Zero-shot	0.614	0.734	0.614	0.635	0.517	0.582	0.464	0.476
	Few-shot	0.690	0.717	0.690	0.692	0.519	0.560	0.436	0.443
	Finetuning	0.765	0.789	0.765	0.758	0.690	0.686	0.679	0.675
InternVL3-38B	Zero-shot	0.784	0.797	0.784	0.785	0.649	0.665	0.603	0.608
	Few-shot	0.806	0.817	0.806	0.807	0.673	0.666	0.630	0.630
	Finetuning	0.767	0.795	0.767	0.758	0.636	0.540	0.535	0.528

Table 2: Classification Macro Metrics for Primary Category and Subcategory by Model and Training Method. Highest values in each column are in **green**, second-highest in **cyan**. P_m, R_m, and F1_m refer to macro precision, macro recall, and macro F1 scores.

vised finetuning delivers substantial improvements across subcategory tasks, underscoring that direct model adaptation is the most effective strategy for specialized harmful content detection.

5.6 Per-Class Detection Performance

Per-class performance analysis reveals systematic variation in detection difficulty that validates the importance of expert clinical knowledge (Table 3 for Task 1 and Table 12, Appendix E.2 for Task 2).

Content Explicitness Hierarchy A clear hierarchy emerges: *Supplement Abuse* achieves the highest performance across all models (up to 0.928 F1) as harmful indicators like steroid mentions (“tren”, “stacking”) provide unambiguous lexical and visual signals. In contrast, *Relationship to Body* and *Relationship to Exercise* categories prove most challenging (F1 \approx 0.70–0.76). These categories require distinguishing legitimate fitness content from harmful messaging along subjective boundaries, creating a dependency on expert knowledge that even large commercial models struggle to navigate without specific tuning.

Cultural Nuance and Data Density Commercial models outperform open-source alternatives on culturally nuanced categories like *Relationship to Masculinity*, likely due to broader pre-training on

sociological concepts. However, expert-labeled finetuning allows open-source models to yield results competitive with commercial systems in specific domains (e.g., InternVL3-38B matches Claude-Sonnet-4 on *Supplements*). Notably, finetuning effectiveness correlates directly with annotation density: while these models excel on data-rich *Irrelevant* content ($N = 548$), they struggle significantly on sparse categories like *Masculinity* ($N = 164$).

5.7 Severity Score Prediction

To assess the ability of VLMs to predict harm severity, we conducted zero-shot prompting experiments across multiple models. Results shown in Table 4 reveal that models typically misestimate severity by approximately half a scale point, indicating reasonable but imperfect alignment with expert clinical judgment. Notably, severity prediction proves more challenging than categorical classification, with all models showing only moderate correlation with ground truth ratings. While these findings demonstrate baseline competence in severity estimation, substantial improvement is needed for reliable moderation deployment.

5.8 Error Analysis

Cross-model error analysis exposes fundamental challenges in characterizing pro-bigorexia content,

Primary Category	Commercial Models						Open-Source Models					
	Claude-4		Gemini-2.5		GPT-4.1		InternVL3-38B			Qwen2.5VL-32B		
	ZS	FS	ZS	FS	ZS	FS	ZS	FS	FT	ZS	FS	FT
Body	0.749	0.757	0.716	0.697	0.740	0.739	0.700	0.733	0.702	0.586	0.693	0.689
Food	0.866	0.856	0.829	0.809	0.807	0.840	0.832	0.842	0.845	0.780	0.854	0.829
Exercise	0.762	0.727	0.734	0.694	0.713	0.728	0.708	0.725	0.719	0.655	0.728	0.720
Supplements	0.922	0.928	0.892	0.882	0.897	0.924	0.898	0.928	0.890	0.870	0.882	0.901
Masculinity	0.867	0.869	0.828	0.771	0.832	0.847	0.793	0.818	0.549	0.735	0.844	0.654
Irrelevant	0.794	0.769	0.826	0.796	0.766	0.796	0.777	0.798	0.843	0.827	0.841	0.863

Table 3: Per-Class F1 Scores for Primary Category Classification by Model and Training Paradigm. ZS = Zero-Shot, FS = Few-Shot, FT = Finetuned. Highest scores per category in **green**, 2nd highest in **blue**.

Model	Severity	
	MAE ↓	ρ ↑
Claude-Sonnet-4	0.679	0.675
GPT-4.1	0.690	0.690
Gemini-2.5-Flash	0.693	0.691
InternVL3-8B	0.805	0.474
InternVL3-38B	0.701	0.607
Qwen2.5-VL-7B	0.794	0.484
Qwen2.5-VL-32B	0.688	0.601

Table 4: Zero-shot Severity Prediction Metrics. MAE: Mean Absolute Error (lower is better); ρ : Spearman’s Rank Correlation (higher is better). Highest values in **green**, 2nd highest in **blue**.

particularly when visual signals conflict with textual context.

Qualitative Case Study We illustrate these challenges through a representative error in an example video (Video_ID 030). The ground truth label is *Relationship to Masculinity* because the on-screen text explicitly ties specific body metrics (“170-190 lbs”, “1.6 shoulder-to-waist ratio”) to female desirability (“what women find”). While Claude-Sonnet-4 (zero-shot) correctly identifies this Masculinity context, both Gemini-2.5-Flash and GPT-4.1 (zero-shot) misclassify it as *Relationship to Body*. This error suggests that the misclassifying models overweight the visual modality, which features prominent shirtless posing, while underweighting the crucial textual overlay that reframes the physique display as a prerequisite for male social worth. This exemplifies the difficulty of distinguishing between pure body objectification and masculinity-driven dysmorphia without deep multimodal reasoning.

Systematic Misclassification Broader quantitative analysis mirrors these qualitative findings.

Models consistently misclassify exercise-related content as *Irrelevant* (0.061–0.143) or confuse *Relationship to Body* with *Relationship to Exercise* (0.102–0.173), reflecting annotator disagreements. Videos demonstrating workout sessions simultaneously touch on exercise, body display, dieting, lifestyle, and motivations that single-label classification cannot capture. Classification consistently confuses *Hormone Therapy* and *Anabolic Steroids* (0.118–0.176 bidirectional misclassification), a clinically critical distinction. Models also frequently misclassify *Supplement Abuse* as *Relationship to Body* (0.082–0.224), suggesting difficulty in recognizing harmful supplement messaging. This likely stems from creators using coded language (e.g., “tren”, “stack”), requiring domain expertise that general-purpose models lack. Additionally, models show systematic bias toward predicting *Muscularity Self-objectification* in subtype classification, with excessive exercise and toxic motivation content frequently misclassified into this category (0.147–0.471). These systematic errors highlight a fundamental challenge: effective pro-bigorexia detection demands not just multimodal capabilities, but clinical and social knowledge to navigate the blurred boundaries between fitness content and promoting unhealthy behaviors.

5.9 Ablation Study

5.9.1 Input Features: Text vs Video

Table 5 reports results of ablation study to assess the contribution of each modality to classification performance on Task 1 with zero-shot prompting. We evaluate GPT-4.1 and InternVL3-38B (image-based) and Gemini-2.5-Flash (video-native). Multimodal fusion yields the best results, with GPT-4.1 and InternVL3-38B achieving F1 scores of 0.792 and 0.785, respectively. Video alone provides

Model	Modality	P _m	R _m	F1 _m
Gemini-2.5-Flash	Audio	0.700	0.423	0.424
	Audio All	0.669	0.471	0.478
	Caption	0.732	0.707	0.709
	OCR	0.717	0.661	0.664
	Text	0.733	0.719	0.719
	Video	0.786	0.765	0.766
	All	0.782	0.776	0.775
GPT-4.1	Audio	0.752	0.444	0.453
	Audio All	0.724	0.488	0.505
	Caption	0.728	0.672	0.679
	OCR	0.741	0.663	0.672
	Text	0.749	0.675	0.681
	Video	0.755	0.697	0.701
	All	0.808	0.792	0.792
InternVL3-38B	Audio	0.632	0.405	0.416
	Audio All	0.630	0.454	0.470
	Caption	0.669	0.590	0.587
	OCR	0.649	0.614	0.618
	Text	0.691	0.630	0.635
	Video	0.697	0.655	0.660
	All	0.797	0.784	0.785

Table 5: Ablation results for task 1 using a single modality input features (text or video). *Audio* modality refers to the audio transcript of the video whose audio is classified as speech; *Audio All* modality refers to the audio transcript of all videos; *Caption* refers to the video description; *OCR* is text within images, and *Text* modality refers to *Caption* + *OCR*. Highest values in each column are in **green**, second-best in **cyan**. P_m, R_m, and F1_m refer to macro precision, macro recall, and macro F1 scores.

the strongest unimodal performance (F1 = 0.660–0.701), highlighting the discriminative power of visual behavioral cues. In contrast, audio consistently underperforms (F1 < 0.453), likely due to background music and variable audio quality. Gemini-2.5-Flash shows narrower modality gaps, suggesting native video processing reduces the need for explicit fusion. Overall, results underscore the value of combining visual and textual inputs for robust pro-bigorexia detection.

5.9.2 Number of Frames per Video

We examine whether providing more video content as input features to the models helps improve results. Our frame-count ablation (Table 6) reveals contrasting patterns between models. Performance of GPT-4.1 shows substantial degradation with increased frames (F1: 0.827→0.792→0.685), while InternVL3-38B maintains stable performance across all densities (F1: 0.785–0.790). This suggests that GPT-4.1 suffers from information overload when processing dense temporal sequences, whereas InternVL3-38B effectively cap-

Model	Frames #	P _m	R _m	F1 _m
GPT-4.1	4	0.832	0.829	0.827
	16	0.817	0.793	0.792
	32	0.739	0.657	0.685
InternVL3-38B	4	0.7967	0.7840	0.7849
	16	0.7988	0.7891	0.7896
	32	0.7922	0.7840	0.7844

Table 6: Frame-count ablation for GPT-4.1 and InternVL3-38B. P_m, R_m, and F1_m are macro precision, macro recall, and macro F1. Highest values in each column are in **green**, second-best in **cyan**.

tures the relatively static visual elements in TikTok pro-bigorexia content. The stability across frame counts indicates that complementary text and audio modalities provide sufficient dynamic contextual information, validating our cost-efficient 4-frame approach.

6 Conclusion

We introduce BIGTOKDETECT to bridge the gap between generic content moderation and the clinical nuances of muscle dysmorphia. While focused on this specific pathology, our work targets a broader class of multimodal mental health risks that evade traditional guardrails. Crucially, our results demonstrate that identifying these subtle behavioral harms relies less on model size and more on domain-specific adaptation; supplying high-quality, expert-annotated data for finetuning proves more effective than relying on massive generalist models. Future work may extend this via RLHF to integrate annotator disagreement as meaningful signal, preserving model utility while capturing the ambiguity inherent in clinical assessment that rigid training often discards. This functionality is particularly critical as short-form video platforms like TikTok increasingly shape youth culture and to combat the use of generative AI in health misinformation. Our framework provides reproducible methodology for the dynamic, behavior-aware moderation systems necessary to address the growing complexity of modern online harms.

Limitations

Dataset and Sampling Constraints Our dataset is limited to English-only TikTok content from 2019–2025, potentially missing cultural variations and platform-specific differences. We focused on English to ensure annotation quality and TikTok as the dominant youth platform. Keyword-

based sampling may overlook subtle or emerging pro-bigorexia content that avoids our taxonomy terms, though we exhaustively developed keywords through previous literature review and expert consultation to maximize coverage.

We did not track creator gender due to the privacy and ethical risks of visual inference. Since pro-bigorexia content typically skews male, while our “Irrelevant” control is likely gender-diverse, this demographic discrepancy presents a potential confounder for future work to address via privacy-preserving methods.

Annotation and Taxonomy Limitations Forced reduction from multi-faceted expert annotations to single labels loses nuanced information. We simplified to single labels due to current LLM limitations in reliable multi-label classification. Our 16 annotators are predominantly female (13/16), potentially introducing gender perspective bias in male-centric bigorexia evaluation; however, this field remains understudied, and we could not find sufficient male experts specializing in bigorexia. Our taxonomy represents a living document that may miss emerging patterns, though we strived for exhaustiveness through iterative expert refinement.

Model Selection Limitations While we selected current leading VLMs across both commercial and open-source categories, the rapidly evolving landscape means we may be missing other capable models. We prioritized models with proven multimodal video understanding capabilities, but acknowledge that newer or specialized architectures might offer different performance characteristics.

Experimental Reproducibility Due to computational cost constraints, we conduct single runs for each experimental configuration without multiple trials or statistical significance testing. While our results establish baseline performance across models and tasks, future work should include multiple experimental runs with statistical analysis to provide more robust performance comparisons and confidence intervals for the reported metrics.

Ethics Statement

This research addresses harmful content classification, raising important ethical considerations. To protect annotators reviewing potentially disturbing pro-bigorexia content, we recruited clinical experts, licensed psychologists, psychiatrists, and doctoral candidates specializing in eating disorders, plus

computational social scientists with extensive experience in body image research. All annotators participated as voluntary research collaborators, fully informed of the project scope, with explicit rights to withdraw without penalty.

To strictly protect user privacy, we maintain strict protocols. While annotators accessed original content to ensure accurate clinical coding, we implemented a comprehensive manual de-identification protocol for the public data release. Four authors conducted a frame-by-frame review of the entire dataset. For every video, we annotated the specific on-screen coordinates of embedded usernames (@<username>), which can shift across frames, and redacted these regions using a white overlay box. Furthermore, we excised all platform-generated outro segments that display creator handles. This manual redaction, combined with the removal of user IDs and metadata, ensures that the stored videos and extracted frames contain no visible creator usernames or TikTok account identifiers. The full dataset will be made available only upon eligible research requests with ethics approval and signed data use agreements.

While supporting mental health research, we acknowledge risks of censorship against legitimate fitness communities. We focus on detection research rather than deployment recommendations. Our work is intended for research purposes and should not be used for clinical diagnosis without proper clinical oversight. The study protocol was reviewed and approved by the University of Southern California Institutional Review Board (IRB application number UP-23-00750).

We acknowledge the assistance of an AI language model in writing and editing portions of this manuscript. All AI-generated content was reviewed and edited by authors. Research design, data collection, analysis, and primary contributions remain entirely human work.

Acknowledgments

We are deeply grateful to our clinical collaborators at the University of California, Los Angeles Health, including Krista Tabuenca, Artha Gillis, Colleen McCord, Brandy Saccacio, Amanda Velkova, Chessa Kabling, Jenny Nguyen, and Danielle Crowe, for generously volunteering their time and expertise to annotate this sensitive content, without whose clinical insights this work would not have been possible

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, and 8 others. 2022. [Flamingo: a visual language model for few-shot learning](#). *Preprint*, arXiv:2204.14198.
- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders (DSM-5®)*, 5th edition. Diagnostic and Statistical Manual of Mental Disorders. American Psychiatric Publishing, Arlington, VA.
- Anthropic PBC. 2025. [System card: Claude opus 4 & claude sonnet 4](#). Technical report, Anthropic PBC. Accessed: 2025-07-27.
- Corey H Basch, Lorie Donelle, Joseph Fera, and Christie Jaime. 2022. [Deconstructing tiktok videos on mental health: Cross-sectional, descriptive content analysis](#). *JMIR Form Res*, 6(5):e38340.
- Anne E. Becker. 2004. [Television, disordered eating, and young women in fiji: Negotiating body image and identity during rapid social change](#). *Culture, Medicine and Psychiatry*, 28(4):533–559.
- Charles Bickham, Bryan Ramirez-Gonzalez, Minh Duc Chu, Kristina Lerman, and Emilio Ferrara. 2025. [Edtok: A dataset for eating disorder content on tiktok](#). *Preprint*, arXiv:2505.02250.
- Guy Cafri, Roberto Olivardia, and Joel Thompson. 2008. [Symptom characteristics and psychiatric comorbidity among males with muscle dysmorphia](#). *Comprehensive psychiatry*, 49:374–9.
- Stevie Chancellor and Munmun Choudhury. 2020. [Methods in predictive techniques for mental health status on social media: a critical review](#). *NPJ Digital Medicine*, 3(1):43.
- Stevie Chancellor, Yannis Kalantidis, Jessica A. Pater, Munmun De Choudhury, and David A. Shamma. 2017. [Multimodal classification of moderated online pro-eating disorder content](#). In *Proceedings of the 2017 ACM CHI Conference on Human Factors in Computing Systems*, pages 3213–3226. ACM.
- Minh Duc Chu, Cinthia Sánchez, Zihao He, Rebecca Dorn, Stuart B. Murray, and Kristina Lerman. 2024. [Leveraging machine learning to identify gendered stereotypes and body image concerns on diet and fitness online forums](#). *Preprint*, arXiv:2407.03551.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3290 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Marita Cooper, Kamryn T. Eddy, Jennifer J. Thomas, Debra L. Franko, Bradley Carron-Arthur, Ani C. Keshishian, and Kathleen M. Griffiths. 2020. [Muscle dysmorphia: A systematic and meta-analytic review of the literature to assess diagnostic validity](#). *International Journal of Eating Disorders*, 53(10):1583–1604.
- Melissa Donati, Ludovica Polidori, Paola Vernillo, and Gloria Gagliardi. 2023. [Building a corpus on eating disorders from TikTok: Challenges and opportunities](#). In *Proceedings of the 9th Italian Conference on Computational Linguistics (CLiC-it 2023)*, pages 515–520, Venice, Italy. CEUR Workshop Proceedings.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Caifeng Shan, and 2 others. 2025. [Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis](#). *Preprint*, arXiv:2405.21075.
- Kyle T. Ganson, Alexander Testa, Rachel F. Rodgers, and Jason M. Nagata. 2025. [Associations between muscularity-oriented social media content and muscle dysmorphia among boys and men](#). *Body Image*, 53:101903.
- Tarleton Gillespie. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press, New Haven, Connecticut.
- David Gimeno-Gómez, Ana-Maria Bucur, Adrian Cosma, Carlos-David Martínez-Hinarejos, and Paolo Rosso. 2024. [Reading between the frames: Multimodal depression detection in videos from non-verbal cues](#). *Preprint*, arXiv:2401.02746.
- Robert Gorwa, Reuben Binns, and Christian Katzenbach. 2020. [Algorithmic content moderation: Technical and political challenges in the automation of platform governance](#). *Big Data & Society*, 7:205395171989794.
- Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, R. Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron J. Weiss, and Kevin Wilson. 2017. [Cnn architectures for large-scale audio classification](#). *Preprint*, arXiv:1609.09430.
- Nick Ara Kamkari. 2025. [The manosphere and men’s health: unpacking the links between online communities, body dysmorphia and erectile dysfunction](#). *Journal of Men’s Health*, 21(4):123–125.

- Douwe Kiela, Suvrat Bhooshan, Hamed Firooz, Ethan Perez, and Davide Testuggine. 2020a. [Supervised multimodal bitransformers for classifying images and text](#). *Preprint*, arXiv:1909.02950.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020b. [The hateful memes challenge: Detecting hate speech in multimodal memes](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 2611–2624. Curran Associates, Inc.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Preprint*, arXiv:2309.06180.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). *Preprint*, arXiv:2301.12597.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Tackling multi-image, video, and 3d in large multimodal models](#). In *arXiv preprint arXiv:2407.07895*.
- Valerie Lookingbill, Ehsan Mohammadi, and Yizhou Cai. 2023. [Assessment of accuracy, user engagement, and themes of eating disorder content in social media short videos](#). *JAMA Network Open*, 6(4):e238897.
- Marcus Ma, Georgios Chochlakis, Niyantha Maruthu Pandiyan, Jesse Thomason, and Shrikanth Narayanan. 2025. [Large language models do multi-label classification differently](#). *Preprint*, arXiv:2505.17510.
- Donald R. McCreary and Doris K. Sasse. 2000. [Drive for muscularity scale \(dms\)](#). Database record.
- Ghofrane Merhbene, Alexandre Puttick, and Mascha Kurpicz-Briki. 2024. [Investigating machine learning and natural language processing techniques applied for detecting eating disorders: a systematic literature review](#). *Frontiers in Psychiatry*, 15:1319522.
- Maria Minadeo and Lizzy Pope. 2022. [Weight-normative messaging predominates on tiktok—a qualitative content analysis](#). *PLOS ONE*, 17(11):e0267997.
- Deborah Mitchison, Jonathan Mond, Scott Griffiths, Phillipa Hay, Jason M. Nagata, Kay Bussey, Nora Trompeter, Alexandra Lonergan, and Stuart B. Murray. 2022. [Prevalence of muscle dysmorphia in adolescents: findings from the everybody study](#). *Psychological Medicine*, 52(14):3142–3149.
- Stuart Murray, Leila Hazery, Tori Shen, Tom Wooldridge, and Jonathan Mond. 2015. [Go big or go home: A thematic content analysis of pro-muscularity websites](#). *Body image*, 16:17–20.
- Stuart B Murray, Jason M Nagata, Scott Griffiths, Jerel P Calzo, Tiffany A Brown, Deborah Mitchison, Aaron J Blashill, and Jonathan M Mond. 2017. [The enigma of male eating disorders: A critical review and synthesis](#). *Clinical Psychology Review*, 57:1–11.
- OpenAI. 2023. [Gpt-4 technical report](#). Online. <https://openai.com/research/gpt-4>.
- OpenAI. 2025. [Gpt-4.1](#). <https://openai.com/index/gpt-4-1/>. Accessed: 2025-07-27.
- Harrison G. Jr. Pope, Anthony J. Gruber, Paulette Choi, Roberto Olivardia, and Katharine A. Phillips. 1997. [Muscle dysmorphia: An underrecognized form of body dysmorphic disorder](#). *Psychosomatics*, 38(6):548–557. Nov–Dec.
- Qwen Team. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115. ArXiv:2412.15115v2.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *Preprint*, arXiv:2212.04356.
- TikTok. 2025. [Query videos – tiktok research api](#). Online; accessed July 2025.
- TikTok Creative Center. [Top hashtags overview](#). <https://ads.tiktok.com/business/creativecenter/hashtag/pc/en>. Accessed January 2025.
- Tao Wang, Markus Brede, Antonella Ianni, and Emmanouil Mentzakis. 2017. [Detecting and characterizing eating-disorder communities on social media](#). In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 91–100.
- Megan Willett. 2023. [Body dysmorphia in boys and men can fuel muscle obsession, doctors say](#). *The Washington Post*.
- Almut Zeeck, Viola Welter, Hasan Alatas, Tom Hildebrandt, Claas Lahmann, and Armin Hartmann. 2018. [Muscle dysmorphic disorder inventory \(mddi\): Validation of a german version with a focus on gender](#). *PLoS ONE*, 13(11):e0207535.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). *Preprint*, arXiv:2403.13372.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Yuchen Duan, Hao Tian, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Yanan Liu, Tan Jiang, Jiapeng Luo, Yi Wang, Yanan He, Tan Lu, Lewei Lu, and 4 others. 2025. [Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models](#). *arXiv preprint arXiv:2504.10479*.

Appendix

A Taxonomy

A.1 Categories Definition

A detailed definition of the categories in our taxonomy is provided in Table 7. These categories are carefully curated and iteratively refined in collaboration with clinical psychologists and domain researchers to ensure both validity and practicality.

Throughout taxonomy development, clinical experts engage in extensive deliberation over edge cases that reveal the central challenge of drawing the line between harmful and legitimate fitness content—a distinction requiring specialized clinical expertise, real-world patient experience, and collaborative consensus-building. For instance, the *Maladaptive Coping* subcategory sparks debate about whether content promoting exercise as emotional regulation constitutes harm, with experts weighing clinical distinctions between healthy stress management and problematic avoidance behaviors. Similarly, *Legal APEDs* discussions center on sarcastic supplement content, where experts consider whether adult humor might inadvertently normalize excessive consumption among adolescent audiences. These deliberations require experts to balance contextual factors, including target audience, creator intent versus potential impact, and clinical thresholds between enthusiasm and pathology. Through structured consensus sessions examining ambiguous cases, annotators establish decision rules that balance clinical rigor with practical annotation consistency.

B Annotation

B.1 Annotator Profiles

We enlist 16 subject-matter experts, ranging from clinical psychologists and social workers to computational social scientists, for video annotation. Our annotators are research collaborators and co-authors who volunteered their clinical expertise with full knowledge of the project’s scope and scientific objectives, ensuring informed participation in this sensitive content annotation task. Table 13 details each annotator’s ID, batch assignment, area of expertise, and gender.

B.2 Annotation Instructions

Annotators received comprehensive guidelines emphasizing safety protocols and decision-making frameworks for handling potentially disturbing

content. Key instructions included: (1) **Content Safety**: annotators were advised to work at their own pace, take regular breaks, and prioritize mental well-being when reviewing harmful content; (2) **Decision Framework**: when encountering ambiguous cases, annotators were instructed to select “Unsure, need consultation” and provide explanatory notes, with textual information (captions, on-screen text) prioritized over conflicting visual content; (3) **Consistency Maintenance**: annotators were asked to reflect on their annotation patterns periodically to avoid drift and maintain consistency across the 300-video batches; (4) **Expert Support**: availability of research coordinators via dedicated communication channels for immediate consultation on challenging cases.

Annotators assess harm severity based on message intensity, including explicit display of muscular physiques, toxic motivational pressure toward unattainable muscularity, rigid and unrealistic dietary demands, endorsement of high-risk behaviors such as steroid use or dangerous workout practices, and the degree of pathological behavioral patterns exhibited in the content.

These protocols ensured both annotator safety and annotation quality while acknowledging the inherently difficult nature of distinguishing harmful pro-bigorexia content from benign fitness material.

B.3 Annotation Platform

The HTML interface (Figure 2) presents each TikTok video with its caption and hashtags (if any). The form beneath the video prompts the annotator to mark if the video is: (1) *Showing One’s Body* (yes/no); includes (2) *Commercial Sponsorship* (yes/no); (3) 2 boxes for *Type of harm* (1 mandatory, 1 optional), consisting of primary categories and subcategories, plus “Relevant but Not Listed,” “Irrelevant,” and “Unsure—Needs Consultation” options; and (4) *Severity of Harm* on a discrete scale ranging from 1 (Not Harmful) to 5 (Extremely Harmful). A text box allows annotators to write additional comments and edge-case observations.

B.4 Annotation Processing

After Round 2, we apply a five-step pipeline to prepare the data for consensus review and final curation:

1. **Annotation Completeness**: verify that each video has exactly two independent annotations.

Primary	Secondary	Definition	Keywords
Relationship to Body	Muscularity Self-Objectification	Idealized imagery emphasizing muscular aesthetics as the primary source of value.	shredded, swole, mensphysique, bodybuilding
	Leanness Self-Objectification	Self-evaluation against a lean, low-fat, highly-toned ideal promoted online.	flat tummy, skinny men physique, small waist fitness
	Muscle Dissatisfaction	Expressing perceived insufficient muscularity despite fitness, fueling negative self-view.	muscles never big enough, muscles not big enough, muscle dysmorphia, bdd men
Relationship to Food	Rigid Food Rules	Obsessive macro/micronutrient tracking and extreme bulking or cutting diets.	aggressive cut, aggressive bulk, shredded diet, macro tracking
	Unsafe Food	Promotion of raw or unsafe foods believed to enhance muscle growth.	liver king diet, raw meat diet for gym, dog food to gain muscles
	Cheat Meals	Large “reward” meals after restrictive dieting that reinforce binge–compensation cycles.	cheatday food, cheatmeal
Relationship to Exercise	Excessive Exercise	Obsessive routines exceeding healthy limits despite injury or life interference.	no rest day, david goggins mentality, train until failure, push your limit
	Predebiting	Treating exercise as punishment or permission to eat, creating guilt cycles.	exercise so I can eat, workout so I can eat, earn your food, train so I can eat
	Maladaptive Coping	Using exercise as the sole coping mechanism to avoid emotional distress.	gym therapy, workout breakup, workout heartbreak, gym mental health, gym fixes everything
	Exercise-Induced Functional Impairment	Prioritizing exercise over essential duties, harming daily functioning.	gym over everything, gym or nothing, gym over friends, skip school for gym
	Toxic Motivation	Demeaning communication that pressures unrealistic fitness standards via shaming or slurs.	gym masculinity, aggressive gym motivation, they don’t know me son, get your ass to the gym
Supplements	Anabolic Steroids	Normalization or endorsement of anabolic-androgenic steroid use with downplayed risks.	tren, anabolic stack, anabolic gear, steroids
	Legal APEDs	Overuse of legal supplements (e.g., creatine, protein, preworkout) beyond recommended doses.	protein powder, whey, creatine, pre workout
	Hormone Therapy	Downplaying risks and spreading misinformation about testosterone replacement therapy.	TRT for gains, testosterone, hormoneboost
Relationship to Masculinity		Links muscle-building and exercise performance to male identity, sexuality, and self-worth.	gym masculinity, be a man gym, embrace masculinity, alpha male gym
Irrelevant		General fitness or lifestyle content unrelated to pro–bigorexia harm.	fyp, tiktok, foryoupage, viral, funny, duet, trending, love, meme, followme, repost, new, awesome, music, cute, video, foryou, fun, diy, ootd, family, lifehack, photography, usa, college, travel, christmas, sport, party, popular, clip, movie, star, moment, tiktokviral, tiktokfamous, tiktokmusic, likeforfollow, recipe, quotes, TikTokChallenge, memories

Table 7: Primary and secondary categories of the Harm Taxonomy for Pro-Bigorexia Content on TikTok, with definitions and search keywords.

EDTok Bigorexia Harm Annotation

If you need more guidance, please refer to the EDTok_Bigorexia Harm Taxonomy [here](#).



\$(video_description)

Showing One's Body

Does this video show any visuals of a person's body? No

Commercial Sponsorship

Is this video commercially sponsored? No

Type of Harm

Select a first main category of harm. Relationship to Body

Select an appropriate subcategory of harm. Muscularity Self-objectification

Select a secondary main category of harm. Relationship to Food

Select an appropriate subcategory of harm. Unsafe Food

Severity of Harm

How much would this content reinforce and glorify ED to a viewer when exposed to this content?



Note

Any additional comments or observations here...

Submit Annotation

Figure 2: Annotation Interface on Amazon Mechanical Turk platform.

- Data Filtering:** remove videos that are unavailable (deleted) or non-English, as indicated in annotators' notes.
- Reliability Assessment:** compute Cohen's κ for both strong (exact primary–secondary match) and weak (primary-only match) agreement to quantify inter-rater reliability.
- Category Refinement Candidates:** collect all videos flagged as “Unsure—Needs Consultation” or “Relevant but Not Listed.” For these, we consult with the annotators to either reassign them to existing categories, extend the taxonomy to cover frequently emerging content types, or mark them as irrelevant.
- Consensus Review and Re-annotation:** filter out all items by agreement level (strong, weak, disagreement), sample a subset of weak and disagreement cases for each annotator pair, hold a batch-specific meeting to discuss and jointly re-annotate selected cases to reach consensus, then return the remaining weak and disagreement items to annotators for in-

dependent re-annotation.

This pipeline ensures data consistency, highlights contentious cases for Round 3 consensus sessions, and informs any necessary taxonomy updates.

B.5 Annotation Statistics

Inter-annotator agreement substantially improves from Round 2 to Final (Tables 3 and 4), with Cohen's κ values increasing from moderate agreement (0.43-0.69 strong κ , 0.59-0.83 weak κ) to good-to-excellent agreement (0.58-0.81 strong κ , 0.78-0.94 weak κ). The consistently high weak κ values (>0.9 in most batches) in the Final round indicate that annotators achieved near-perfect agreement on the broader annotation categories.

Batch	Round 2		Final	
	Strong κ	Weak κ	Strong κ	Weak κ
1	0.551	0.613	0.616	0.775
2	0.570	0.588	0.762	0.841
3	0.510	0.612	0.680	0.815
4	0.551	0.785	0.811	0.937
5	0.563	0.685	0.787	0.924
6	0.555	0.723	0.742	0.938
7	0.430	0.739	0.577	0.931
8	0.685	0.825	0.769	0.938

Table 8: Cohen's κ coefficients for inter-annotator agreement across annotation rounds. **Strong** κ denotes exact agreement on the Subcategory label; **Weak** κ denotes agreement on the Primary Category only. Reliability improved substantially in the Final round following consensus adjudication.

C Data

We collected videos through the official TikTok Research API after submitting a research proposal that was reviewed and approved by TikTok. This API accesses publicly posted content where users have consented to public visibility through TikTok's terms of service. We implemented additional privacy protections by anonymizing user identifiers and exposing only video content and captions to annotators.

Our dataset comprises TikTok videos annotated for hierarchical classification across two levels of granularity. The primary category classification categorizes content into six broad categories related to body image and health behaviors, while the subcategory classification provides fine-grained classification into 16 specific subcategories. The severity score scale defines the intensity of the harm emotion. The severity scores are most heavily con-

Primary Category	Count
Relationship to Body	1442
Relationship to Exercise	1233
Relationship to Food	824
Supplement Abuse	558
Relevant but Not Harmful	411
Relationship to Masculinity	406
Irrelevant	401

(a) Primary Categories

Subcategory	Count
Muscularity Self-objectification	905
Leanness Self-objectification	238
Muscle Dissatisfaction	240
Rigid Food Rules	375
Cheat Meals	193
Excessive Exercise	379
Maladaptive Coping	213
Toxic Motivation	308
Anabolic Steroids	195
Other (please specify in Note)	217

(b) Subcategories

Table 9: Distribution of raw individual annotations (before consensus aggregation). Since each video received two independent annotations, these counts reflect the total number of labels assigned by annotators, whereas Table 1 reflects the final unique video counts after aggregation.

centrated in the 1.0 to 1.5 range (Figure 5), indicating that lower severity levels are the most common. The frequency then gradually decreases as scores increase, with relatively few cases exceeding 4.0. This suggests that mild severity is predominant in the dataset. Figure 7 shows that most content falls in the moderate-harm range, though *Supplement Abuse* and *Relationship to Food* score higher on average, likely because markers such as extreme dieting practices or steroid use are more readily identifiable to annotators.

We remove corrupted files (e.g., missing metadata) and non-English videos during preprocessing. The dataset is split into train/test sets, maintaining approximately a 3:1 ratio with stratified sampling based on task labels and downsampling of dominant categories to improve balance.

As shown in Table 1, the primary category task contains 1,966 training and 588 test samples. Training data exhibits natural class imbalance reflecting real-world distributions, ranging from 489 samples (“Relationship to Body”) to 66 samples (“Relationship to Masculinity”), while the test set maintains a strictly balanced representation (98 samples per

category) for fair evaluation. The subcategory task operates on a filtered subset of 1,472 training and 462 test samples. Training samples range from 200 (“Irrelevant”) to 28 (“Predebtting Exercise”), while test data maintains a balanced distribution for major categories (34 samples each) with reduced representation for rare categories (4–10 samples).

Zero-shot and few-shot approaches are evaluated solely on test sets, while finetuning models are trained on the respective task’s training set and evaluated on the corresponding test set.

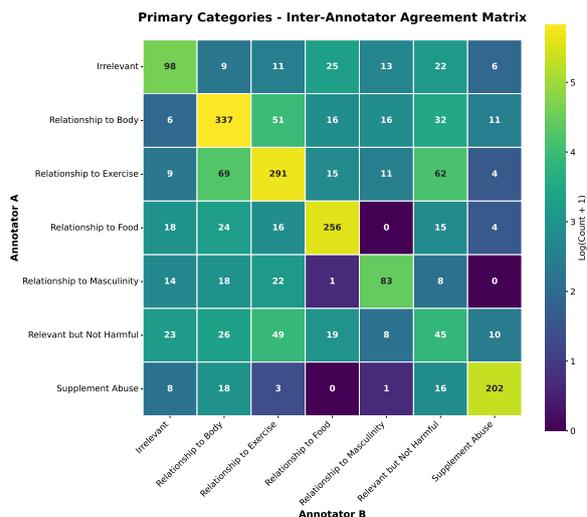


Figure 3: Inter-annotator agreement matrix for primary categories. Each cell shows the number of videos where Annotator A assigned the row category and Annotator B assigned the column category. Diagonal elements represent perfect agreement, while off-diagonal elements indicate disagreements between annotators.

D Modeling

D.1 Experiment Setup

We implement the VLMs using Transformers-based (for InternVL3) and vLLM (Kwon et al., 2023) (for Qwen-2.5VL) implementations through the LlamaFactory framework (Zheng et al., 2024), which provides efficient inference and serving capabilities for VLMs. We used pre-trained VLMs according to their respective licenses and terms of use: commercial API-based models (GPT-4.1, Claude-Sonnet-4, Gemini-2.5-Flash) under their standard API terms, and open-source models (InternVL3, Qwen2.5-VL) under their permissive licenses for research use.

Due to hardware constraints, we exclude InternVL-38B and InternVL-78VL from certain experiments. Our current infrastructure is incompati-

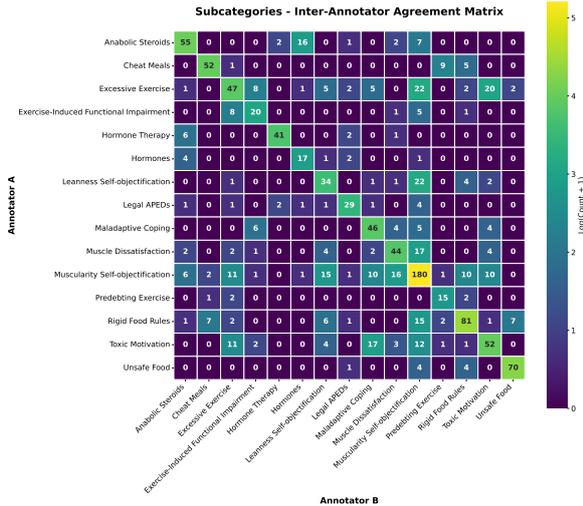


Figure 4: Inter-annotator agreement matrix for subcategories. Each cell shows the number of videos where Annotator A assigned the row subcategory and Annotator B assigned the column subcategory. The matrix reveals patterns of confusion between semantically related subcategories and overall annotation consistency.

ble with the vLLM version required for multi-node deployment of these larger models. Since these models exceed the memory capacity available on single nodes in our cluster, we cannot accommodate their substantial memory requirements within our computational environment.

We consider conducting experiments with LLaVA-NeXT-Video (Liu et al., 2024), which achieved SOTA open-source performance on Video-MME (Fu et al., 2025). However, we exclude LLaVA-NeXT-Video from our study since its context window limitation (4,096 tokens) cannot accommodate our prompting setup (text + video/frames), which often exceeds that limit.

All zero-shot experiments on open-source VLMs are conducted on $8 \times$ NVIDIA H100 GPUs. For few-shot and finetuning experiments, we use $32 \times$ NVIDIA A100 GPUs for Qwen2.5-VL-7B, Qwen2.5-VL-32B, InternVL3-8B, and InternVL3-38B, and $64 \times$ NVIDIA A100 GPUs for Qwen2.5-VL-72B.

We focus on evaluating two tasks: primary category and subcategory classification. For each task, we split the data into training and test sets at a 3:1 ratio using stratified sampling. The test set is strictly balanced across all classes, while the training set is adjusted by downsampling the majority categories to mitigate class imbalance. Detailed split statistics and sampling procedures are provided in Appendix C.

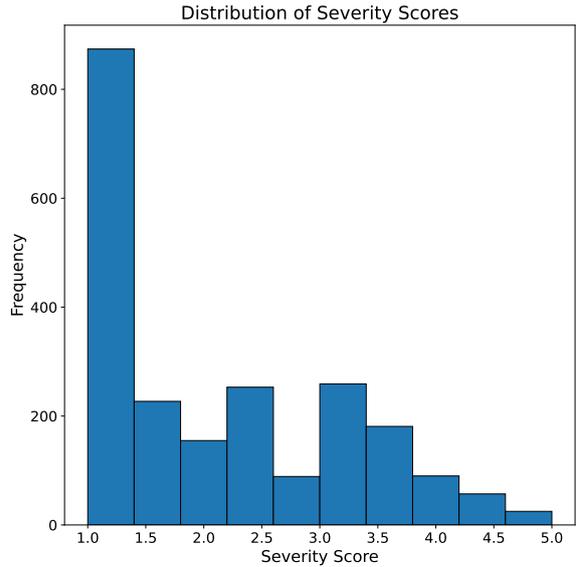


Figure 5: Distribution of Severity Scores. The histogram shows most values concentrated between 1.0 and 2.0, with fewer cases at higher severity levels.

Model	Version	Size (B)	Type
GPT-4o	2024-08-06	—	API
Claude Sonnet 4	2024-02-24	—	API
Gemini 2.5 Pro	2025-06-17	—	API
InternVL3	2024-10-04	8, 38	Open
Qwen2.5-VL	2025-06-05	7, 32, 72	Open

Table 10: Model version dates, parameter counts, and types (API vs. Open-source) for video-based VLMs.

D.2 Prompt Templates

We use a JSON-based prompt structure following the standard role-content format for chat-based LLMs called ‘sharegpt’. Each prompt consists of:

- A sequence of images from the video (sampled frames)
- A user message containing the video caption, embedded text and audio transcription
- A single assistant response with the predicted label

An example format of prompt used for zero-shot inference/supervised finetuning and few-shot inference is shown in Figures 8 and 9 and Figure 10 and 11.

E Results

E.1 Hyperparameter Robustness Check

To validate our choice of using default temperature settings ($T = 1.0$) for commercial API models,

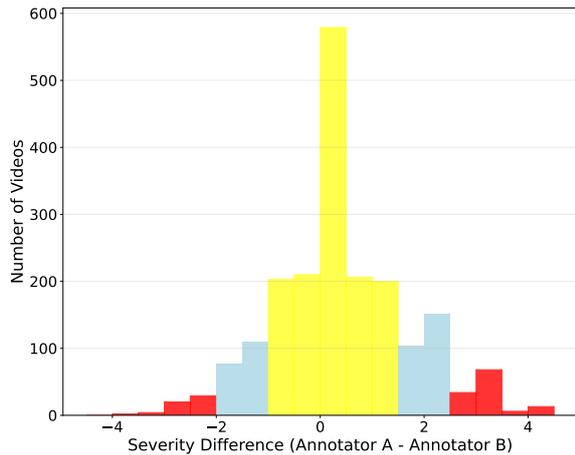


Figure 6: Histogram of severity score differences between annotators. Positive values indicate higher severity by Annotator A. Colors denote agreement levels: green (exact), yellow (± 1 point), red (> 2 points).

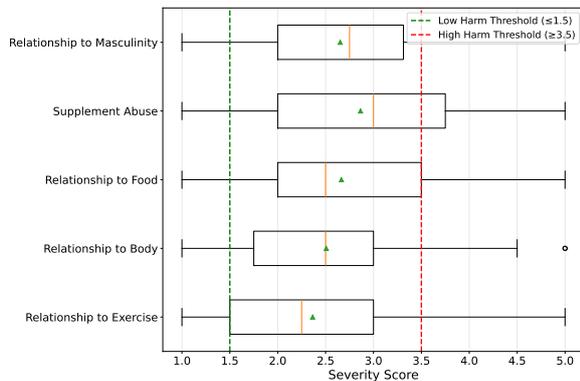


Figure 7: Box plot of severity scores across the primary types. Each box shows the distribution of annotated severity scores, with mean markers included. Dashed vertical lines indicate the harm thresholds, distinguishing low-harm content (≤ 1.5) from high-harm content (≥ 3.5).

we re-evaluated Task 1 (Primary Category Classification) using a deterministic setting ($T = 0.1$). Table 11 compares the performance across both settings.

The results demonstrate that performance remains highly stable, with F1 score variations confined within a narrow range (± 0.01). For instance, Claude-Sonnet-4 shows a slight improvement at $T = 0.1$ (+0.008 F1), while GPT-4.1 shows a minor decrease (-0.004 F1). These marginal differences indicate that the specific temperature choice for high-performing commercial models does not significantly alter the comparative conclusions drawn in the main paper.

Model	Temp	Acc.	P_m	R_m	$F1_m$
Claude-Sonnet-4	1.0	0.829	0.832	0.829	0.827
Claude-Sonnet-4	0.1	0.837	0.841	0.837	0.835
GPT-4.1	1.0	0.796	0.808	0.792	0.792
GPT-4.1	0.1	0.787	0.809	0.787	0.788
Gemini-2.5-Flash	1.0	0.805	0.807	0.805	0.805
Gemini-2.5-Flash	0.1	0.810	0.813	0.810	0.809

Table 11: Robustness check for commercial API models on Task 1 (Primary Category) comparing default ($T = 1.0$) vs. deterministic ($T = 0.1$) temperature settings. P_m , R_m , and $F1_m$ denote macro precision, recall, and F1.

E.2 Per-Subtype Detection Analysis

Table 12 details the F1 scores for all 15 subcategories, revealing significant performance variance driven by content explicitness.

Explicitness Hierarchy Subcategories characterized by distinct lexical keywords or unambiguous visual objects achieve the highest detection rates. *Cheat Meals* (Claude-4 ZS: 0.899) and *Supplement Abuse* subtypes (e.g., *Anabolic Steroids*, *Hormone Therapy*) consistently achieve F1 scores above 0.80, benefiting from clear indicators like specific drug names (e.g., “tren”, “TRT”). Conversely, performance drops sharply for ambiguous categories requiring subjective behavioral interpretation, such as *Muscularity Self-objectification* and *Excessive Exercise* (< 0.60 F1), reflecting the difficulty in distinguishing harmful behaviors from normative training without deep context.

Model-Specific Strengths Commercial models demonstrate superior reasoning capabilities on psychologically nuanced categories, with Claude-4 (Few-Shot) achieving the highest score on *Maladaptive Coping* (0.787). In contrast, open-source models like InternVL3-38B excel in categories with strong visual or explicit signals, leading performance in *Unsafe Food* and *Anabolic Steroids* (Few-Shot: 0.831), confirming their utility for detecting overtly harmful content.

Subcategory	Commercial Models						Open-Source Models					
	Claude-4		Gemini-2.5	GPT-4.1		InternVL3-38B			Qwen2.5VL-32B			
	ZS	FS	ZS	ZS	FS	ZS	FS	FT	ZS	FS	FT	
Muscularity Self-objectification	0.451	0.431	0.512	—	0.562	0.495	0.578	0.384	0.349	0.538	0.391	
Leanness Self-objectification	0.676	0.635	0.645	—	0.706	0.636	0.590	0.500	0.533	0.625	0.561	
Muscle Dissatisfaction	0.655	0.738	0.774	0.714	0.765	0.677	0.765	0.654	0.676	0.754	0.806	
Rigid Food Rules	0.721	0.698	0.714	0.642	0.618	0.769	0.750	0.703	0.607	0.725	0.730	
Unsafe Food	0.812	0.848	0.833	0.800	0.812	0.889	0.857	0.899	0.691	0.759	0.879	
Cheat Meals	0.899	0.781	0.866	0.742	0.750	0.836	0.889	0.700	0.712	0.901	0.818	
Anabolic Steroids	0.706	0.735	0.750	0.765	0.730	0.781	0.831	0.781	0.754	0.825	0.690	
Legal APEDs	0.700	0.818	0.667	0.818	0.800	0.833	0.696	0.667	0.556	0.636	0.696	
Hormone Therapy	0.769	0.794	0.712	0.788	0.824	0.783	0.794	0.778	0.776	0.794	0.800	
Excessive Exercise	0.449	0.510	0.408	0.507	0.538	0.433	0.500	0.600	0.267	0.459	0.476	
Maladaptive Coping	0.712	0.787	0.655	0.560	0.759	0.489	0.471	0.607	0.419	0.490	0.635	
Exercise-induced Functional Impairment	0.571	0.533	0.500	—	0.500	0.308	0.533	0.429	0.333	0.333	0.353	
Toxic Motivation	0.607	0.552	0.603	0.571	0.594	0.563	0.523	0.448	0.533	0.508	0.516	
Relationship to Masculinity	0.779	0.740	0.621	0.806	0.712	0.618	0.691	0.691	0.588	0.759	0.649	
Irrelevant	0.513	0.470	0.629	0.475	0.532	0.559	0.568	0.658	0.527	0.545	0.743	

Table 12: Per-Class F1 Scores for Subcategory Classification by Model and Training Paradigm. ZS = Zero-Shot, FS = Few-Shot, FT = FineTuned. Highest scores per subcategory in **green**, second-highest in **cyan**. Dashes indicate missing data.

ID	Batch	Expertise	Gender
A1	1	clinical psychology doctoral candidate specializing in male eating disorders and muscle dysmorphia intervention research.	Male
A2	1	senior researcher in computational social science focused on social media dynamics and their implications for mental health.	Female
A3	2	health communication professor using computational social network methods to study online health behaviors.	Female
A4	2	medical student focusing on clinical and translational research in body image disorders and digital health interventions.	Female
A5	3	clinical psychology doctoral candidate specializing in digital mental health and disordered eating prevention.	Female
A6	3	communication doctoral candidate investigating how interactive media-driven social comparisons affect body image outcomes.	Female
A7	4	clinical professor and psychiatrist studying the psychopathology and treatment of eating disorders and muscle dysmorphia.	Male
A8	4	clinical professor and pediatric psychologist specializing in adolescent body image disorders.	Female
A9	5	child and adolescent psychiatrist specializing in developmental psychopathology and body image disorders.	Female
A10	5	clinical social worker specializing in child and adolescent body image disorders and psychosocial interventions in neuropsychiatric care.	Female
A11	6	occupational therapist specializing in functional rehabilitation and psychosocial support for eating-disorder patients in a neuropsychiatric hospital setting.	Female
A12	6	clinical dietitian specializing in nutritional assessment and dietary management for eating and body-image disorders.	Female
A13	7	mental health nurse practitioner specializing in the assessment and treatment of child and adolescent psychiatric disorders.	Female
A14	7	board-certified child and adolescent psychiatrist specializing in mood and behavioral disorders.	Female
A15	8	clinical social worker with inpatient psychiatric experience in providing psychotherapy to adult patients.	Female
A16	8	doctoral candidate specializing in detecting disordered behaviors in social media.	Male

Table 13: Annotator assignments, areas of expertise, and gender.

VLM Zero-Shot/SFT Prompt - Task 1: Primary Category Classification

<video>You are an expert classifier of harmful pro-muscle-dysmorphia content on TikTok.

Taxonomy:

Type 1: Relationship to Body:

- Subtype 1: Muscularity Self-Objectification: Idealized imagery emphasizing muscular aesthetics.
- Subtype 2: Leanness Self-Objectification: Self-evaluation against a lean ideal.
- Subtype 3: Muscle Dissatisfaction: Perceived insufficient muscularity.

Type 2: Relationship to Food:

- Subtype 1: Rigid Food Rules: Obsessive tracking, bulking, or cutting.
- Subtype 2: Unsafe Food: Promoting unconventional muscle-enhancing foods.
- Subtype 3: Cheat Meals: Binge-compensation cycles.

Type 3: Supplement Abuse:

- Subtype 1: Anabolic Steroids: Downplaying risks of steroid use.
- Subtype 2: Legal APEDs: Excessive legal supplement usage.
- Subtype 3: Hormone Therapy: Promoting TRT as performance enhancer.

Type 4: Relationship to Exercise:

- Subtype 1: Excessive Exercise despite harm.
- Subtype 2: Predebting: Exercising to justify eating.
- Subtype 3: Maladaptive Coping: Sole coping method.
- Subtype 4: Functional Impairment.
- Subtype 5: Toxic Motivation.

Type 5: Relationship to Masculinity:

- Subtype 1: Linking muscles to male identity and worth.

Type 6: Irrelevant:

- Subtype 1: Content unrelated to the above.

Carefully consider all sources of information about this video:

VIDEO_ID: <Anonymized Video ID>

Caption: <TikTok Video Caption>

Audio transcript: <Audio Transcriptions using Whisper>

Embedded text: <OCR text from video>

###TASK###

Classify the video into one of the following types:

Relationship to Body, Relationship to Food, Supplement Abuse, Relationship to Exercise, Relationship to Masculinity, or Irrelevant.

Only output the type label, no explanations or subtypes.

Figure 8: Prompt used for zero-shot inference and supervised finetuning for Task 1.

VLM Zero-Shot/SFT Prompt - Task 2: Subcategory Classification

<video>You are an expert classifier of harmful pro-muscle-dysmorphia content on TikTok.
Taxonomy:
Type 1: Relationship to Body:
• Subtype 1: Muscularity Self-Objectification: Idealized imagery emphasizing muscular aesthetics.
• Subtype 2: Leanness Self-Objectification: Self-evaluation against a lean ideal.
• Subtype 3: Muscle Dissatisfaction: Perceived insufficient muscularity.

Type 2: Relationship to Food:
• Subtype 1: Rigid Food Rules: Obsessive tracking, bulking, or cutting.
• Subtype 2: Unsafe Food: Promoting unconventional muscle-enhancing foods.
• Subtype 3: Cheat Meals: Binge-compensation cycles.

Type 3: Supplement Abuse:
• Subtype 1: Anabolic Steroids: Downplaying risks of steroid use.
• Subtype 2: Legal APEDs: Excessive legal supplement usage.
• Subtype 3: Hormone Therapy: Promoting TRT as performance enhancer.

Type 4: Relationship to Exercise:
• Subtype 1: Excessive Exercise despite harm.
• Subtype 2: Predebting: Exercising to justify eating.
• Subtype 3: Maladaptive Coping: Sole coping method.
• Subtype 4: Functional Impairment.
• Subtype 5: Toxic Motivation.

Type 5: Relationship to Masculinity:
• Subtype 1: Linking muscles to male identity and worth.

Type 6: Irrelevant:
• Subtype 1: Content unrelated to the above.

Carefully consider all sources of information about this video:

VIDEO_ID: <Anonymized Video ID>
Caption: <TikTok Video Caption>
Audio transcript: <Audio Transcriptions using Whisper>
Embedded text: <OCR text from video>

###TASK###
Classify the video into one of the following subtypes: Muscularity Self-Objectification, Leanness Self-Objectification, Muscle Dissatisfaction, Rigid Food Rules, Unsafe Food, Cheat Meals, Anabolic Steroids, Legal APEDs, Hormone Therapy, Excessive Exercise, Predebting, Maladaptive Coping, Exercise-Induced Functional Impairment, Toxic Motivation, Relationship to Masculinity, or Irrelevant.

Only output the specific subtype label, no explanations or other text.
If none apply, output Irrelevant.
Use the exact subtype names, not the type names.

Figure 9: Prompt used for zero-shot inference and supervised finetuning for Task 2.

VLM Few-Shot Prompt. Task 1: Primary Category Classification

You are an expert classifier of harmful pro-muscle-dysmorphia content on TikTok.

Taxonomy:

Type 1: Relationship to Body:

- Subtype 1: Muscularity Self-Objectification: Idealized imagery emphasizing muscular aesthetics as the primary source of value.
- Subtype 2: Leanness Self-Objectification: Self-evaluation against a lean, low-fat, highly-toned ideal.
- Subtype 3: Muscle Dissatisfaction: Expressing perceived insufficient muscularity despite having a muscular physique.

Type 2: Relationship to Food:

- Subtype 1: Rigid Food Rules: Obsessive macro/micronutrient tracking and restrictive dietary practices to rapidly gain muscle mass (bulking) or lose fat (cutting).
- Subtype 2: Unsafe Food: Promotion of unconventional foods believed to enhance muscle growth.
- Subtype 3: Cheat Meals: Large "reward" meals after restrictive dieting that reinforce binge-compensation cycles.

Type 3: Supplement Abuse:

- Subtype 1: Anabolic Steroids: Normalization or endorsement of anabolic-androgenic steroid use with downplayed risks.
- Subtype 2: Legal APEDs: Overuse of legal supplements (e.g., creatine, protein, pre-workout) beyond recommended doses.
- Subtype 3: Hormone Therapy: Downplaying risks and spreading misinformation about testosterone replacement therapy, promoted as a performance enhancer or "anti-aging" treatment without proper medical diagnosis.

Type 4: Relationship to Exercise:

- Subtype 1: Excessive Exercise: Extreme exercise routines exceeding healthy limits despite injury or life interference.
- Subtype 2: Predebting: Treating exercise as punishment or permission to eat.
- Subtype 3: Maladaptive Coping: Using exercise as the sole coping mechanism to avoid emotional distress.
- Subtype 4: Exercise-Induced Functional Impairment: Prioritizing exercise over essential duties, harming daily functioning.
- Subtype 5: Toxic Motivation: Demeaning communication that pressures unrealistic fitness standards via shaming or slurs.

Type 5: Relationship to Masculinity:

- Subtype 1: Linking muscle-building and exercise performance to male identity, sexuality, and self-worth.

Type 6: Irrelevant:

- Subtype 1: Content without muscle-obsession, restrictive diets, supplement/AAS promotion, or extreme exercise. Includes dance trends, memes, travel/cooking vlogs, general wellness (e.g., yoga), and pure entertainment unrelated to body-image narratives.

Here are some examples from the training data:

Example 1:

<video>

Caption: <Video Caption>

Audio transcript: <Audio Transcription from Whisper>

Embedded text: <OCR Text>

Classification: <Type>

.

.

.

Example 12:

<video>

Caption: <Video Caption>

Audio transcript: <Audio Transcription from Whisper>

Embedded text: <OCR Text>

Classification: <Type>

<video>

VIDEO_ID:

Caption:

Audio transcript:

Embedded text:

###TASK###

Classify the video into one of the following types: Relationship to Body, Relationship to Food, Supplement Abuse, Relationship to Exercise, Relationship to Masculinity, or Irrelevant.

Only output the type label, no explanations or other text. If none apply, output Irrelevant.

Don't use a subtype, such as "Muscularity Self-Objectification" or "Unsafe Food".

Valid outputs: Relationship to Body, Relationship to Food, Supplement Abuse, Relationship to Exercise, Relationship to Masculinity, Irrelevant

Figure 10: Prompt used for few-shot inference for Task 1.

VLM Few-Shot Prompt. Task 2: Subcategory Classification

You are an expert classifier of harmful pro-muscle-dysmorphia content on TikTok.

Taxonomy:

Type 1: Relationship to Body:

- Subtype 1: Muscularity Self-Objectification: Idealized imagery emphasizing muscular aesthetics as the primary source of value.
- Subtype 2: Leanness Self-Objectification: Self-evaluation against a lean, low-fat, highly-toned ideal.
- Subtype 3: Muscle Dissatisfaction: Expressing perceived insufficient muscularity despite having a muscular physique.

Type 2: Relationship to Food:

- Subtype 1: Rigid Food Rules: Obsessive macro/micronutrient tracking and restrictive dietary practices to rapidly gain muscle mass (bulking) or lose fat (cutting).
- Subtype 2: Unsafe Food: Promotion of unconventional foods believed to enhance muscle growth.
- Subtype 3: Cheat Meals: Large "reward" meals after restrictive dieting that reinforce binge-compensation cycles.

Type 3: Supplement Abuse:

- Subtype 1: Anabolic Steroids: Normalization or endorsement of anabolic-androgenic steroid use with downplayed risks.
- Subtype 2: Legal APEDs: Overuse of legal supplements (e.g., creatine, protein, pre-workout) beyond recommended doses.
- Subtype 3: Hormone Therapy: Downplaying risks and spreading misinformation about testosterone replacement therapy, promoted as a performance enhancer or "anti-aging" treatment without proper medical diagnosis.

Type 4: Relationship to Exercise:

- Subtype 1: Excessive Exercise: Extreme exercise routines exceeding healthy limits despite injury or life interference.
- Subtype 2: Predebting: Treating exercise as punishment or permission to eat.
- Subtype 3: Maladaptive Coping: Using exercise as the sole coping mechanism to avoid emotional distress.
- Subtype 4: Exercise-Induced Functional Impairment: Prioritizing exercise over essential duties, harming daily functioning.
- Subtype 5: Toxic Motivation: Demeaning communication that pressures unrealistic fitness standards via shaming or slurs.

Type 5: Relationship to Masculinity:

- Subtype 1: Linking muscle-building and exercise performance to male identity, sexuality, and self-worth.

Type 6: Irrelevant:

- Subtype 1: Content without muscle-obsession, restrictive diets, supplement/AAS promotion, or extreme exercise. Includes dance trends, memes, travel/cooking vlogs, general wellness (e.g., yoga), and pure entertainment unrelated to body-image narratives.

Here are some examples from the training data:

Example 1:

<video>

Caption: <Video Caption>

Audio transcript: <Audio Transcription from Whisper>

Embedded text: <OCR Text>

Classification: <Type>

.

.

Example 12:

<video>

Caption: <Video Caption>

Audio transcript: <Audio Transcription from Whisper>

Embedded text: <OCR Text>

Classification: <Type>

<video>

VIDEO_ID:

Caption:

Audio transcript:

Embedded text:

###TASK###

Classify the video into one of the following subtypes: Muscularity Self-Objectification, Leanness Self-Objectification, Muscle Dissatisfaction, Rigid Food Rules, Unsafe Food, Cheat Meals, Anabolic Steroids, Legal APEDs, Hormone Therapy, Excessive Exercise, Predebting, Maladaptive Coping, Exercise-Induced Functional Impairment, Toxic Motivation, Relationship to Masculinity, or Irrelevant.

Figure 11: Prompt used for few-shot inference for Task 2.

Video ID	Subcategory	Video Description	Embedded Text	Audio Transcript
001	Relationship to Body	my lil frame	me being a 5'8 male but with lean body	and my little frame and my sweet little girl voice. It exudes something in people that...
002	Relationship to Masculinity	#menshealth #menshealthtips #cynic	@<redacted>	The reason that there is such a void of male role models for young men is because we literally do not have adult males in our population anymore. What do I mean by that? There's no adult males. Like, there's grown men everywhere, but they're not actually physiologically adult human males...
003	Supplement Abuse	this is totally healthy right? #recipes #creatine	@<redacted>	So I told a co-worker I was just taking my creatine in water, just water, because it was unflavored. And they told me you're supposed to take it with like a meal or like sugar so then you can absorb it better...
004	Relationship to Food	<redacted> dinner!	@<redacted>	Pulled beef with hard shelled eggs Mmm It's crazy. It actually tastes like the delicious pulled beef That's in with the garlic oh and the onion and the lime juice booyah, yeah Sweet potatoes, the purple variety...
005	Relationship to Exercise	I am girl	MAKING THE GYM YOUR WHOLE PERSONALITY no train legs	girls who make going to the gym their entire personality it is just so crazy to me I genuinely do not know how you guys do it I could never be one of those girls anyway someone I know just asked me what day it was and I replied with leg day
006	Irrelevant	Two in #shorts #gymshorts	@<redacted>	ugly, ugly, ugly aura. Bro, if your shorts, especially if you're hitting leg day, are below your knees, you're genuinely cooked, bro. These, I got them here on the TikTok shop. This is a size medium. I'm five foot 10, 165 pounds, and these things fit so good...

Table 14: Example entries for Type classification showing captions and TikTok metadata. The actual videos can be found in the data attachment in the submission platform.

Video ID	Subcategory	Video Description	Embedded Text	Audio Transcript
007	Muscularity Self-objectification	training hard or hardly training ?	sounds of my workout on push day	Yeah, we need some weight, should be a good one today. Bet you gotta hit me fast, you know old...
008	Leanness Self-objectification	my lil frame	me being a 5'8 male but with lean body	and my little frame and my sweet little girl voice. It exudes something in people that...
009	Muscle Dissatisfaction	#gymtok	When u start working out to look better but now u have body dysmorphia	fucking go let's go I guess
010	Rigid Food Rules	I gained 35lbs in 90 days #lifting #gym #bulking	35lbs in 90 days I've considered ever since 7th grade but change	I gained 35 pounds in 90 days. I've considered myself a runner ever since 7th grade...
011	Unsafe Food	Liver King Dinner to start the week	<redacted>	Mmm, steak and potatoes, sweet potato fritters, Liver King Chef Lionel, that's how he calls them. So good. We got...
012	Cheat Meals	NO MORE.. #gym #cheatday #gains	us after cheat day turned into cheat month	No more cheese fillets. No more McDonald's. No more chicken wings. No more chicken snobby with some sour. No more...
013	Excessive Exercise	Gym rats! #gym #gymtok #mothersoftiktok #thick #newyearsresolutions #muscles #gains #proteinshake #lawenforcement	It's time for us normal gym peeps to get our routines back!!!	Listen, I got my ass the gym today and I didn't want to go. So did you go today? Well...
014	Predebtng Exercise	You CANNOT our train a bad diet #diet #fitness #exercise	I CAN EAT WHAT-EVER I WANT AS LONG AS I EXERCISE YUP YOU HEARD ME WHILE EXERCISE IS IMPORTANT..	I can eat whatever I want as long as I exercise. You said what? Yup, you heard me. While exercise...
015	Maladaptive Coping	Gym Therapy #gymtherapy #gym #doordie #dobetter #military	@<redacted>	Gym therapy is a term that was coined by Shredded Vets that describes the psychological effects of working out consistently...
016	Exercise-Induced Functional Impairment	@<redacted> - I would skip school just to go to the gym. I was going to do whatever it took.	I would skip school to go to the gym.	lucky enough like when I was in year 9 and year 10 the Knights were giving me a day off...
017	Toxic Motivation	put in the effort #gym #motivation #fitness-motivation #workout	TD PRESSURE	It's fucking different. Gotta really come kill me. I'm built different. I train different. I work different. I am different...
018	Anabolic Steroids	#<redacted>	POV: Steroids gives you acne HAVE FUCKING I MAN LIKE PIMPLE SHIT COLLAB MIGHT PIMPLE UH ASIAN	dude I think I have like a fucking I got like a cyst man you know like where it's a...
019	Legal APEDs	Sustenanceeee #fyp #xyzbca #gymtok #creatine	POV: you get caught using creatine	Are there any drugs in this house? If there are, you better find them and give them to me immediately...
020	Hormone Therapy	Day 5 on TRT #trt #lift #weightlifting #gym #workout	Day 5 on TRT	What up guys, it's day five. So on day five, the doctor calls you with your full blood work done...

Table 15: Example entries for Subtype classification showing captions and TikTok metadata. The actual videos can be found in the data attachment in the submission platform.