

The First Mass Protest on Threads: Multimodal Mobilization and AI-Generated Visuals in Taiwan’s Bluebird Movement

Tracy Weener¹, Ho-Chun Herbert Chang¹

¹Program in Quantitative Social Science, Dartmouth College, Hanover, NH 03755 USA

Abstract

The 2024 Bluebird Movement in Taiwan marked one of the largest youth-led protests in the country’s democratic history, mobilizing over 100,000 demonstrators in response to parliamentary reforms. Unlike the 2014 Sunflower Movement, Bluebird unfolded within a transformed digital environment dominated by Threads, Meta’s new microblogging platform that draws 24% of its global traffic from Taiwan. Leveraging a dataset of 62,321 posts and 21,572 images, this study analyzes how protest communication developed across textual and visual modalities. We combine LLM zero-shot annotation, gradient-boosting trees, and SHAP explainers, to disambiguate the supply and demand of attention. Results reveal three dynamics: (1) a partisan gap between algorithmic exposure and user engagement, with anti-DPP content surfaced more widely but anti-KMT and pro-DPP content more actively recirculated; (2) textual repertoires centered on commemorations, personal testimonies, and calls to action as key drivers of virality; and (3) a bifurcation in visual strategies, where human photographs concentrated exposure and discussion, while AI-generated animal symbols circulated as mobilization tools and partisan attacks. The exposure-engagement gap demonstrates how Threads functions as both amplifier and filter of democratic contention, suggesting **dual publics** operating under distinct logics. Generative AI further reshapes symbolic repertoires through **kawaii toxicity**, or partisan attacks cloaked in cute aesthetics.

Introduction

In May 2024, more than 100,000 demonstrators took to the streets of Taipei to protest parliamentary reforms that opponents argued undermined checks and balances in Taiwan’s Legislative Yuan (Wang et al. 2024). Triggered by legislative maneuvers expanding investigatory powers in ways that would concentrate authority in the Kuomintang (KMT) and Taiwan People’s Party (TPP) coalition, the protests evolved into the Qing Niao (Blue Bird) Movement, one of the largest recall campaigns in modern democracies, with more than 30 legislators facing potential removal from office. As the only Mandarin-speaking democracy, Taiwan has witnessed significant protests before—notably the 2014 Sunflower Movement—but Bluebird exemplifies critical shifts in how democratic contention unfolds: mass migration to a

new public platform, deployment of AI-generated visuals at scale, and new paradigms in platform-driven censorship.

Shortly before the movement emerged, Meta launched Threads in July 2023, positioning it globally as a competitor to Twitter (now X). While Threads experienced modest adoption in most markets, it saw extraordinary growth in Taiwan. By April 2025, Taiwan accounted for nearly a quarter of Threads’ global web traffic, making it a uniquely important site for public discourse and civic organization (Meta 2023). Threads filled a vacuum that other popular platforms like LINE (closed messaging) and Instagram (network-structure) could not serve. It became the “Twitter Taiwan never had,” distributing attention across multiple demographics and enabling the kind of real-time public discourse that had been absent from Taiwan’s digital ecosystem.

During the Bluebird protests, Threads became the primary space for coordination and debate. But a fundamental question emerges: what happens if a platform’s algorithm privileges content that is different than what protestors want themselves? As early as February 2024, Meta platforms announced it would taper back political content surfaced to its users (Instagram 2024). Knowing this, activists exploited Mandarin’s morpho-syllabic properties, where the name “Qing Niao” bore both visual and auditory similarity to “Qing Dao,” the protest location. This linguistic adaptation circumvented not state censorship but platform moderation. Fundamentally, activists confronted a divergence between what the platform’s algorithm surfaced and what users chose to amplify, which we characterize as an algorithmic exposure and social engagement gap.

Leveraging 62,321 posts and 21,572 images from Threads between May 2024 and June 2025, we analyze how this first massive protest unfolded on Threads. We focus on three interrelated dynamics: the relationship between algorithmic visibility and user-driven circulation, the textual and visual strategies that secure attention, and the role of AI-generated imagery in mobilization and partisan attack. The proliferation of AI-generated animal symbols—blue birds, winter deer, toads, goblins—introduces what we term **cute toxicity** (or **kawaii toxicity**): political attacks cloaked in aesthetics of cuteness. Building on scholarship examining emotional and visual contagion in social movements (Gerbaudo 2012; McAdam, Tarrow, and Tilly 2001), we show how generative AI reshapes symbolic repertoires by packaging partisan

animosity in adorable aesthetics (Allison 2006; Ngai 2012), with implications extending well beyond Taiwan.

Literature Review

Platform-Mediated Protest

Social movement scholarship has long emphasized the role of visibility in mobilization. Classical theories describe cycles of contention—waves of protest that rise, diffuse, and decline as challengers and authorities interact (McAdam, Tarrow, and Tilly 2001). Within these cycles, focusing events trigger moral shock, converting private grievances into public action (Jasper 1998). The underlying assumption is straightforward: if movements can broadcast these events and messages widely enough, they can mobilize support. Visibility meant power.

The advent of social media appeared to amplify this visibility-power relationship. Platforms like Twitter and Facebook lowered broadcasting costs, especially the rapid dissemination of moral-emotional content (Brady et al. 2017; Tufekci 2017). However, platforms are not neutral amplifiers but active curators (Gillespie 2018). Contemporary platforms employ algorithmic ranking systems that selectively surface content based on predicted engagement, user behavior, and platform policy objectives (Bozdog 2013). These systems intervene between what is posted and what is seen, creating a “marketplace of attention” where algorithms work with networks to shape exposure (Webster 2014).

This curation creates a fundamental gap between what is posted (supply), what algorithms show (exposure), and what users circulate (engagement) (Bandy and Diakopoulos 2021). Rather than uniformly amplifying movements, platforms fragment them into overlapping information environments. This divergence is not a new problem. Researchers have distinguished between what elites supply and what publics demand. Munger and Phillips (2022) showed that YouTube’s recommendation algorithm shapes which political content gains exposure (supply-side) versus which content audiences actively seek and engage with (demand-side). Their analysis reveals that supply and demand operate according to different rationales: algorithms optimize for watch time and engagement, while users optimize for identity reinforcement and information seeking.

Importantly, exposure doesn’t simply reflect aggregated user preferences. Gillespie (2018) argues platforms make editorial decisions—what to surface, what to suppress, what to monetize—that shape public discourse. These decisions create an “algorithmic public” that may diverge from the “social public,” or the conversation users would create through their own circulation choices. Asymmetric use also exists in user bases. Freelon, Marwick, and Kreiss (2020) find systematic differences in how political groups engage with platforms: right-leaning activists favor high-volume posting and strategic coordination, while left-leaning activists favor spontaneous individual expression. These asymmetries mean engagement patterns reflect not just message resonance but also group-level strategic differences and platform-specific affordances.

Multimodal Communication and AI-Generated Content

While textual communication has received substantial scholarly attention, recently, scholars have documented the centrality of visual communication in contemporary movements. The 2020 George Floyd protests demonstrated how a single video could trigger global mobilization, functioning as visceral moral shock that text alone could not achieve (Chang, Richardson, and Ferrara 2022). This visual turn reflects platforms’ increasing emphasis on image and video content over text, as well as publics’ preferences for emotionally immediate content (Bonilla and Rosa 2015). Memes, in particular, are critical protest infrastructure, which are defined as units of cultural information that spread through imitation and variation.

Generative AI represents a phase change in symbolic production costs. Text-to-image models enable anyone with a smartphone to produce high-quality images in seconds at close to zero marginal cost. This democratization produces broader participation in aesthetic production and proliferation of symbols that serve coordinating functions (Chang et al. 2026).

A growing body of work examines AI-generated images on dimensions such as hyperbole, realism, and absurdity. A critical aesthetic dimension is *kawaii*—Japanese for “cute” theorized as politically significant. Cute aesthetics in Japanese popular culture enable soft power projection and identity formation. Crucially, cuteness creates complex affective responses—simultaneously evoking vulnerability and manipulation (Ngai 2012) by making exploitation appear benign (Miller 2017). Therefore, partisan animosity and toxicity is an underexplored area in political campaigning and social movements.

Protests in Chinese: Censorship Evasion and Linguistic Innovation

Research on the evasion of censorship often occurs under authoritarian contexts. King, Pan, and Roberts (2013) demonstrate that Chinese censors permit criticism of individual officials but systematically suppress content with collective action potential, leading activists to develop linguistic workarounds—homophonic substitutions, metaphors, and coded language that evade keyword filtering. “Crab culture” on Chinese social media denotes how activists use animal metaphors and visual symbols to discuss taboo topics (Ng 2021). These studies establish that censorship pressure drives linguistic innovation, with activists exploiting language’s flexibility to communicate despite restrictions.

Mandarin presents unique affordances for censorship evasion due to its morpho-syllabic writing system. A single sound can be represented by multiple characters with different meanings, and characters can be visually altered while remaining phonetically recognizable. This linguistic feature enables activists to communicate “in plain sight” while evading keyword filters that platforms deploy.

However, this literature focuses almost exclusively on state censorship in authoritarian regimes. Protestors increasingly face what Roberts (2018) terms “corporate cen-

sorship”—content moderation by private platforms pursuing commercial or reputational objectives, deployed at a transnational level rather than being country-specific. As a democracy, corporate censorship is precisely what Bluebird activists confronted.

Threads and Taiwan

Threads, launched by Meta (the parent company of Facebook and Instagram) on July 5, 2023, is a microblogging platform designed for sharing short-form text, multimedia content, and facilitating public conversation, tightly integrated with Instagram accounts and social graphs. Initially conceived under the internal codename “Project 92,” Threads rapidly positioned itself as a competitor to Twitter—now X—leveraging Meta’s existing infrastructure to gain 10 million users within its first seven hours and over 100 million sign-ups in its first five days, becoming the fastest-growing consumer app in history (Meta 2023; Perelli and Bradley 2025; Zahn 2023). Meta explicitly designed Threads as a platform for more engaging public discourse and creator–fan interactions, leveraging parasocial relationships—one-sided social relationships in which individuals feel a sense of intimacy with a public figure, even though the interaction is not reciprocal (Giles 2002; Horton and Wohl 1956; Tukachinsky and Stever 2018). As of August 2025, Threads reported 400 million users across more than 100 countries (Brandon 2025).

Taiwan accounts for an astounding 23.75% of Threads’ global web traffic, the highest among all countries (AsiaPac 2025). Prior to Threads, Taiwan was dominated by a few social platforms. LINE, a chatroom app, is Taiwan’s most dominant social platform used by 94% of Taiwan’s population. Facebook is used by 17.1 million users as of January 2025, representing 93.6% of the population. Lastly, Instagram serves 11.3 million users in Taiwan, which corresponds to 49.1%. These platforms previously formed the core of Taiwanese politics, with significant levels of discourse in the 2020 and 2024 Taiwanese elections (Chang and Fang 2024; Chang, Haider, and Ferrara 2021).

In Taiwan’s lexicon, Threads quickly earned the humorous nickname “cui” (meaning “crispy” in Mandarin). Notable political figures in Taiwan, including former President Chen Shui-bian, have used Threads in populist, civics-focused ways—like answering students’ legal questions. Due to its public-facing nature, Threads became the Twitter Taiwan never had. Threads became a new gathering space, especially for young users in Taiwan, to discuss political content, and eventually, political mobilization (Yang 2024).

The general election of Taiwan, held in January 2024, featured a three-way horse race between the Democratic Progressive Party (DPP), the Kuomintang (KMT), and Taiwan People’s Party (TPP). Taiwan’s main political cleavages center on identity and territory, unlike traditional left-right divisions, and social movements have directly shaped party formation (Nachman 2025). In the past, presidential elections in Taiwan were a bipartisan race between the DPP and KMT, split between cross-strait attitudes. Growing concern with issues such as energy policy and the domestic economy allowed third parties like the TPP to emerge (Weener, Chang,

and Horiuchi 2025). Although President Lai Ching-te initiated a record third presidency for the DPP, the KMT and TPP held a majority in the Legislative Yuan (LY), equivalent to Congress in the United States. Through a majority alliance, the KMT–TPP coalition could thus effectively block initiatives by the DPP. Many citizens voiced concerns that if legislative voting was done solely on the basis of partisanship, then this LY would diminish Taiwan’s democracy.

The Bluebird Movement was a youth-led civic protest movement in Taiwan that erupted in May 2024 in response to legislative maneuvers by the opposition KMT and TPP (Wang et al. 2024). The protests were triggered by proposed parliamentary reforms, which expanded the legislature’s investigatory powers in ways that could weaken checks and balances and concentrate power in the KMT–TPP coalition. Beginning on May 17, 2024, thousands of students and young people gathered outside the Legislative Yuan in Taipei. By June 4, 2024, the movement raised 2.6 million dollars to display an ad in Times Square (Everington 2024).

As the Bluebird protests grew, activists called not only for the revocation of the reforms, but the recall of specific legislators who had championed them. This was seen by some as a constitutional remedy: if lawmakers had acted in ways that undermined checks and balances, the people could exercise democratic oversight through recalls. Building on the movement’s energy, civic groups—and later the DPP—launched what became known as the “Great Recall” campaign. This unprecedented effort sought the removal of up to 31 legislators, mostly from the Kuomintang (KMT), whom activists accused of obstructing democratic governance.

Research Questions

The Bluebird movement provides a **unique vantage point** for understanding contemporary protest. It captured a rare case where a significant portion of a country migrates to an entirely new platform. In one of the few Mandarin-speaking democracies, the main language provides unique insights as to how linguistic features circumvent not government, but platform-based, censorship. Lastly, it provides insight regarding how AI is visually deployed to mobilize citizens.

Our study thus focuses on the following interrelated research questions:

1. **RQ1:** What is the divergence, if any, between exposure and engagement, and across partisanship?
2. **RQ2:** What textual and visual communication strategies predict engagement?
3. **RQ3:** How is generative AI deployed in protest and counter protest?

Methods

Data Collection

We collected Threads posts related to the Bluebird Movement using the Meta Content Library (MCL) by saving the browsing sessions. We utilized only one keyword, Bluebird (qīngniǎo), and filtered for posts only containing traditional Chinese. This yielded a dataset of 62,321 unique posts from




Figure 1: Annotated timeline of the Bluebird Movement between 05/01/2024 to 06/01/2025.

Threads between 05/01/2024 to 06/01/2025, capturing 13 months of the movement. Figure 1 shows the time-series of the posts during this period. This data included the username, date, post text, and crucially, exposure and engagement metrics such as likes, comments, reposts, and views.

As a quick note on views, according to MCL’s documentation, views are the number of times the post appeared on screen, not counting the owner’s own views. A post has no view count if there were fewer than 100 views. Additionally, research from Meta utilizing similar metrics define a view if a post is surfaced for greater than a certain visual proportion and number of seconds, which is meant to prevent overcounting from scrolling. As a test of quality, a random sample of 100 posts yielded 97% directly related to Taiwan or social mobilization; the other 3% were explicitly about mobilization outside of Taiwan, such as in Spain and France. Coupled with the fact that such a large amount of global traffic emerges from Taiwan and the unique geopolitical distribution of traditional Chinese (used only by Taiwan and Hong Kong, the latter for which no longer has Meta access), the quality of data is extremely high. However, it is unknown if there are foreign actors creating fake accounts; we discuss this as an avenue for additional investigation. Using the MCL urls, we then downloaded images associated with the posts, which yielded 21,572 unique photos. Not all posts have images, and some posts have multiple images.

Data Augmentation with Language Models

We rely on zero-shot learning to annotate posts without task-specific labeled data. In zero-shot setups, a model assigns labels it has never been explicitly trained on by leveraging shared semantic structure or broad pretraining. Early work framed this as mapping inputs and labels into a common semantic/attribute space to enable transfer (Palatucci et al. 2009; Socher et al. 2013; Frome et al. 2013; Xian, Schiele,

and Akata 2017).

In NLP, large pretrained language models exhibit strong zero-shot ability via instruction-following and prompt-based reformulations of classification (e.g., as cloze or NLI tasks), enabling accurate labeling from natural-language task descriptions alone (Radford et al. 2019; Brown et al. 2020; Yin, Hay, and Roth 2019; Gao, Fisch, and Chen 2021; Liu et al. 2023). In our pipeline, we operationalized zero-shot labeling by prompting GPT-4o with a detailed schema and decision rules—without providing examples—and constraining outputs to a predefined JSON format.

Text annotation. For every post, the model extracted basic identifiers (languages, hashtags, locations, music references, dates/times, user tags, emojis). Building on existing typologies of protest communication, we also created a custom macro-set that included the following categories: Condemnation, Support, Personal Testimony, Humor/Satire, Call to Action, News Sharing, Cultural/Historical Context, Question/Information Seeking, Commemorative, Other. The model also flagged likely authorship (human vs. AI) and generated a ≤ 120 -word summary. Tone and style were rated on 1–5 scales for sentiment toward the Bluebird Movement, formality, and aggressiveness. Content coding captured subjects; domestic vs. international politics; references to animals, students, politicians, influencers, and national-flag emojis. Domestic-issue indicators included death penalty, separatism, Legislative Yuan, occupation, recalls, misinformation, LGBTQ+, Indigenous rights, energy, budget, stock prices, electricity prices, overseas Taiwanese, healthcare, and democracy. Party variables recorded mentions of DP-P/KMT/TPP, with 0–100 probabilities of supportive or oppositional stance toward each party. This prompt is provided in Appendix A.

Image annotation. For every image, the model (i) performed basic identification (typology, real vs. AI, artwork/-

cartoon status, short description, realism, meme likelihood, and a higher-level visual category); (ii) extracted any embedded text verbatim (OCR-style) and associated metadata (language(s), hashtags, locations, dates, links/QRs); (iii) coded visual content (objects and probabilities for blue birds, other animals and types, Taiwan/other maps, protests, references to Taiwanese/Hong Kong social movements, human presence, counts, signs, demographic descriptors, politicians, food, sexual content, merchandise, national flags by country, style cues such as kawaii/anime); (iv) coded political content (party references and rally likelihood), in-group/out-group partisan strategies (probabilities of support/attack by party); (v) cross-platform/media references (social platforms, news outlets, outlet language/type); (vi) alliances/international affairs (countries, Taiwan/ROC/China/PRC/CCP, Hong Kong, Milk Tea Alliance, Indo-Pacific/East/Southeast Asia, wars); and (vii) domestic issues (death penalty, Legislative Yuan, occupations, recalls, misinformation, LGBTQ+, Indigenous rights, energy). All probabilistic outputs were constrained to 0–100% and all fields enforced via schema validation. We used deterministic decoding and strict JSON validation (see Appendix A) and validations can be found in Appendix B.

Note on Data Ethics

Access to Meta Content Library requires pre-approval and therefore has a data-sharing agreement. Additionally, while LLM labeling may generate bias, we validated the results manually to ensure suitable accuracy and no imbalance in false negatives and positives. Additionally, prior research has shown labeling elements directly, rather than vague themes, can improve reliability.

Additionally, we were careful that our visualizations only include macro-aggregation of phenomenon, and the presented visuals are ones that appear multiple times. As a broader point, while our re-identification risk is low, replicability does suffer due to our inability to share the data directly. This is a trade-off commonly observed in modern data sharing agreements.

Gradient Boosting Trees and SHAP Explainers

Gradient boosting trees (GBTs) are an ensemble machine learning method that builds a predictive model in a stage-wise manner by sequentially fitting decision trees to the residuals of prior trees. At each iteration, a new tree is trained to approximate the negative gradient of the loss function with respect to the current model's predictions, effectively performing functional gradient descent in function space (Friedman 2001). By iteratively adding these weak learners, GBTs combine the flexibility of decision trees with the power of boosting to capture complex, non-linear relationships while maintaining strong generalization performance. Regularization techniques such as shrinkage (learning rate), subsampling, and tree depth constraints are commonly used to reduce overfitting and improve robustness (Hastie, Tibshirani, and Friedman 2009).

Modern scalable implementations such as XGBoost (Chen and Guestrin 2016), LightGBM (Ke et al. 2017), and CatBoost (Dorogush, Ershov, and Gulin 2018)

extend the original framework with efficient algorithms, parallelization, and categorical feature handling, making GBTs one of the most widely adopted methods for classification, regression, and ranking tasks across domains. This study uses CatBoost, which is well-suited for regressing with categorical variables through combinations of categorical features (Prokhorenkova et al. 2018). We included a 75-25 train-test split, where the model was trained on 75% of the data then validated on a 25% random sample. The learning rate was set to 0.01 with a total of 1,500 epochs, with the tree depth set to 12. This was the result of a hyperparameters grid search over tree depth (6, 9, 12), learning rate (0.03, 0.06), L2 regularization strength (3, 10), and Bayesian bootstrap settings with bagging temperatures of 0 and 2, to provide controlled stochasticity in sampling.

While machine learning techniques yield higher accuracies than canonical statistics, their “black box” nature has limited their interpretability for social science research. In recent years, SHAP explainers have become a useful tool for computing feature importance. Based on Shapley values in game theory, the idea is to calculate the utility contributions of each player to a coalition, based on different power sets of players (Hart 1987). Instead of players, SHAP evaluates power sets of features and their contribution to minimizing error in the model (Lundberg and Lee 2017).

Results

We present results in three parts. First, we document the exposure-engagement gap across partisan frames, revealing how algorithmic visibility and user circulation operate according to distinct logics. Second, we identify textual strategies that predict virality, emphasizing commemorative and testimonial content. Third, we analyze visual communication patterns, with particular attention to AI-generated imagery as a tool for both mobilization and attack.

Textual Analysis

Figure 2a shows the partisan framing of messages in the overall dataset. Analysis of 62,321 posts reveals partisan animosity as the dominant frame, with anti-DPP (14.3%) and anti-KMT (10.6%) content far exceeding supportive messaging. In contrast, only around 7% of posts are pro-DPP and 6% are anti-TPP.

Turning to topics, Figure 2b shows the distribution of key topics in the overall dataset. By far, democracy is the most frequently evoked issue, covering more than one fifth of all content. The Legislative Yuan (LY, 9.8%), references to misinformation (8.6%), and the recall itself (7.4%) take up the rest. These results have some immediate implications. While previous literature has shown that Threads is dominated by young people and DPP supporters, these results show that the supply of DPP-leaning content (i.e., anti-KMT and pro-DPP) is roughly in parity with KMT-leaning content.

This partisan balance in content production makes the exposure-engagement divergence particularly striking, as differences in algorithmic visibility and user circulation cannot be attributed to underlying supply asymmetries—or content that is posted (Zha and Chang 2024). We further consider these stances based on key metrics in exposure and




Figure 2: Textual proportion of Bluebird Threads content, with the (a) partisan framing and (b) topical mentions within all posts.

engagement. Figure 3 shows the top four partisan frames, weighted by their exposure and engagement then normalized against each other. For views, anti-DPP narratives are the most frequent, followed by anti-KMT narratives. However, for reposts, we observe a substantial decrease in anti-DPP reposts, relative to the number of anti-KMT reposts. This suggests that while the feed-ranking algorithm favors anti-DPP content, actual user engagement massively favors sharing anti-KMT content and does not favor anti-DPP content. This is corroborated by the increase in pro-DPP content, exceeding even anti-DPP content for reposts.

These patterns reveal a fundamental asymmetry. Anti-DPP content received 28% more views than anti-KMT messaging despite comparable prevalence (14.3% vs. 10.6% of posts). However, anti-KMT content received 45% more reposts. Pro-DPP frames similarly exceeded their view share in repost and like distributions. For replies, anti-DPP content attracted disproportionate activity, accounting for 34% of all reply volume despite representing just 14% of posts.

This divergence between exposure and engagement metrics reveals patterns in how content circulates. Views reflect what appears in users' feeds, while reposts and likes capture what users choose to amplify. Replies may signal contestation as much as endorsement. This divergence suggests dual publics, one constituted by algorithmic feed-ranking, another by social network circulation. These publics are also dueling, with the algorithm favoring anti-DPP content while users amplify anti-KMT and pro-DPP content. These answer




Figure 3: Top four partisan stances weighted by exposure (views) and engagement (reposts, likes, and replies).

RQ1.

Lastly, to identify the key variables that are driving virality, leveraging SHAP values. Figure 4a) shows the SHAP values for the textual regressor. Being right (left) of the vertical line indicates a positive (negative) effect on likes; red (blue) dots correspond to high (low) levels for each specific variable. All variables along the y-axis are ranked according to their overall importance. The model yielded an R^2 of 0.42. Checks for colinearity (VIF) can be found in the Appendix Table 4.

The most important feature that predicts likes is the category of discussion. The top five category means include posts that are Commemorative (523.93), Personal Testimony (354.00), News Sharing (319.06), Calls to Action (311.26), and Cultural and Historical Context (305.38). Commemorative posts in particular mark protest milestones and reference historical parallels to the Sunflower Movement and invoked collective memory. Other continuous variables show clear patterns: tonally, formal language and negative sentiment lead to more likes. For topic, posts related to political rallies, domestic politics, and the recall yield strong positive coefficients. Lastly, anti-partisan framing also yields the greatest virality.

Visual Analysis

Visual framing follows similar patterns. Figure 4 shows the SHAP values for a CatBoost model built on 93 visual elements. A few clear patterns emerge. Posts that contain actual photographs with textual or visual reference to Taiwan have a significant increase in virality. Similar to the textual case, images with reference to domestic politics yield more attention, whereas ones referencing international politics have generally less. Checks for colinearity (VIF) can be found in the Appendix Table 3. This answers **RQ2**.

Crucially, the visual category is the strongest predictor of likes. To examine this more closely, Figure 5 shows the attention funnel of four visual categories. Overall, images with humans yield the greatest share of views, reposts, and likes, roughly evenly split between crowd-based and non-crowd human photos. This suggests direct references to humans

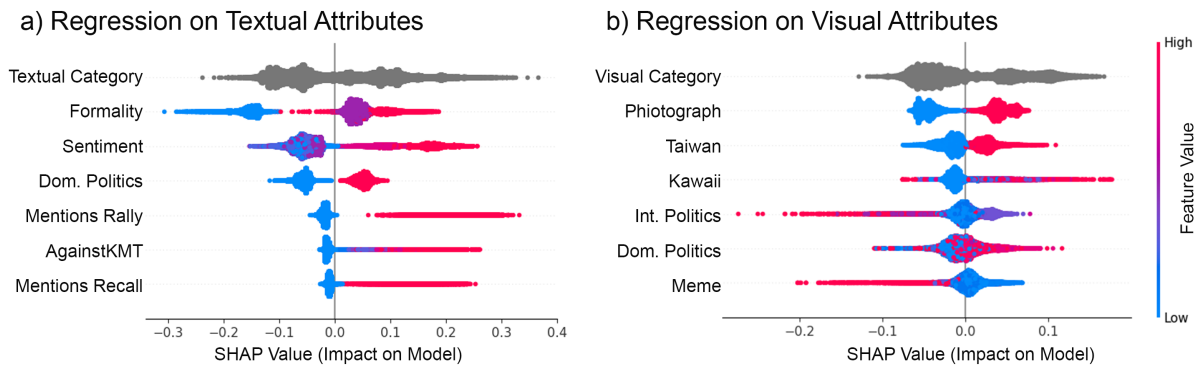


Figure 4: SHAP explainers of the CatBoost Regressor that predicts engagement (likes) based on a) textual and b) visual features.

yield greater exposure and engagement.

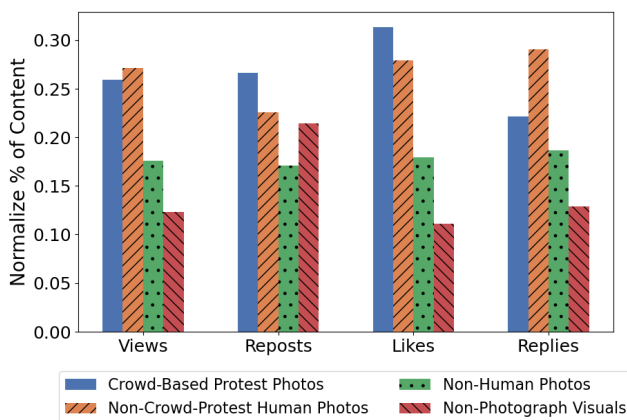


Figure 5: Attention funnel for visual framing categories across exposure and engagement metrics (views, reposts, likes, and replies).

However, a big shift emerges with replies, where non-crowd human photos yield a significant number of replies, whereas crowd-based protest photos yield significantly fewer replies. These human photos are likely pictures of legislators, which prompt fervent discussion, while protest crowd photos are viral but do not elicit discussion. Non-human photos have roughly the same share across the funnel. On the other hand, non-photograph visuals have a substantial boost in reposts, which suggests they are used as visuals for protest mobilization and information sharing.

Diving deeper into the category of non-photograph visuals—content intended to mobilize—one interesting commonality is the large amount of animal and plant related imagery, created using generative AI. Table 1 summarizes the 10 different types of symbols utilized during protest discourse and Figure 6 gives examples of AI-generated images.

Pro-DPP symbols are typically derived via *Xieyin geng*, or puns based on homophones, and explicitly translated to some type of animal. In particular, Blue Bird Winter Deer (qīngniǎo dōnglù) is a homophone for Qing-Dao East Road, the street where the first demonstration of the movement

took place. Whale Island (jīngdǎo) is another homophone that was used, but less popular. Figure 6a shows an AI-generated image of a bird-deer hybrid, while Figure 6b shows a whale encircling Taiwan. These symbols are animals endemic to the country, most exemplified by the Formosan Black Bear. Although the black bear does not have a direct pun, it is often regarded as a symbol of Taiwanese nativism and identity.



Figure 6: Samples of popular AI-generated images in the dataset, including a) blue bird and winter deer, b) whale island, c) a blue bird toad, and toad goblins for d) and e).

However, AI-generated images can also be used to attack. Anti-DPP symbols include goblins and toads, shown in Figures 6c–e. Goblins were initially an anti-TPP term after alleged sexual harassment in the party. However, after DPP whistleblowers accused high-ranking members of the party of sexual harassment, the term evolved to become more associated with anti-DPP sentiment. This association was exacerbated by a television series inspired by the scandal.

The nickname “Pro-Martial Law Toad” originated from a post by the Democratic Progressive Party (DPP) caucus in Taiwan’s Legislative Yuan on social media. In that post, the DPP expressed support for South Korean President Yoon Suk-yeol’s declaration of martial law, describing it as a move

Animal/plant	Translation	Origin	Political Orientation
qīngniǎo	Blue bird	Homonym for Blue Island (qīngdǎo) street, to prevent suppression	Pro-DPP
dōnglù	Winter deer	Homonym for East Road, part of protest location	Pro-DPP
jīngdǎo	Whale	Homonym for Qing-Dao (qīngdǎo) street	Pro-DPP
hēixióng	Black bear	Common symbol for Taiwanese nativism	Pro-DPP
Lányīng	Blue eagle	Refers to KMT supporters (“eagles eating blue birds”)	Pro-KMT
xiǎocǎo	Little grass	Refers to TPP supporters, from popular song	Pro-TPP
Gēbùlín	Goblin	Used in harassment scandals; shifted to anti-DPP	Anti-DPP
chánchú	Toad	Refers to DPP legislators supporting martial law	Anti-DPP
Lài háma	Toad	Derogatory nickname for Lai Ching-te (sound pun)	Anti-DPP
kǎoyā	Roast duck	Linked to Chinese identity; pro-China supporters	Anti-KMT

Table 1: Animal and plant imagery in the Bluebird Protests.

to “safeguard the system of freedom and constitutional governance.” This statement sparked widespread controversy both domestically and internationally. In particular, users on South Korean online forums voiced strong dissatisfaction with the DPP’s stance, even going so far as to call the party a “disgusting toad.” This symbol had additional synergy as President Lai was also frequently called a toad by his political opponents, due to another name for toad sharing a word with his surname (i.e., Lai Hama).

In sum, the case of Taiwan demonstrates how generative AI can be used in modern social movements. Reposted more than it is liked or viewed, it serves as a way to mobilize. More crucially, it exemplifies a new paradigm: hiding toxic content using cute images, as another nefarious consequence of the aesthetics of cuteness (Ngai 2012). Moreover, textual and visual strategies converged on similar principles. High-virality content tends to reinforce collective identity through commemorative narratives and protest photos, or circulated as low-friction mobilization tools through personal testimony and AI-generated symbols. This answers **RQ3**.

Discussion

The Bluebird Movement illustrates how a newly emergent platform can become central to democratic contention. Our analysis of over 62,000 Threads posts reveals that the dynamics of protest communication on this platform cannot be reduced to the volume of content produced. Instead, the results show an asymmetry between what the algorithm amplifies and what citizens choose to endorse. Anti-DPP posts received greater exposure through the feed-ranking system, yet user engagement disproportionately elevated anti-KMT and pro-DPP content (Chang and Fang 2024; Bennett and Segerberg 2012). This suggests that platform logics and social networks may work at cross-purposes: while algorithms surface content that aligns with particular partisan attacks, publics recirculate frames that defend their own side or mobilize against opponents. Such a pattern resonates with recent work emphasizing the dual logics of algorithmic curation and connective action in networked protest (Bennett and Segerberg 2012; Freelon, Marwick, and Kreiss 2020; Tufekci 2017). In this sense, Threads did not merely mirror offline cleavages but actively structured which political identities were rendered visible and which gained traction.

The exposure-engagement gap we document, where anti-DPP content receives 28% more views but 45% fewer reposts than anti-KMT content, points to the possibility of overlapping publics operating according to distinct logics (Bandy and Diakopoulos 2021). Views reflect what appears in users’ feeds; reposts and likes capture what users choose to amplify. Replies may signal contestation as much as endorsement. Rather than a unified public sphere, these divergent metrics suggest parallel information environments where algorithmic curation and user intent do not align uniformly. **Dual publics** describes a pattern where platform visibility and user circulation produce parallel rather than unified information environments.

Equally important is the role of textual repertoires. Posts that commemorated events, shared personal testimony, or called others to action were the strongest predictors of virality. Rather than short bursts of humor or casual commentary, it was messages embedded in collective memory, affect, and civic duty that traveled farthest. This aligns with longstanding theories of moral shock and collective identity in social movements (McAdam, Tarrow, and Tilly 2001; Jasper 1998; Gerbaudo 2012) and with empirical findings that moral-emotional language fuels diffusion on social media (Brady et al. 2017; Chang, Richardson, and Ferrara 2022). Even within a platform optimized for brevity and casual conversation, protest actors crafted texts that linked personal voice to collective stakes. Formal tone and negative sentiment further amplified engagement, suggesting that moments of political threat are framed most effectively through seriousness rather than irony, as a contrast to prior accounts that emphasize humor in online activism (Penney 2017).

Visual communication deepened these dynamics in distinctive ways. Human photographs commanded the largest reach and sparked conversation, particularly when they depicted legislators. These images personalized accountability, directing attention toward elites and inviting direct confrontation in replies. By contrast, AI-generated imagery functioned less as a mode of deliberation and more as a vehicle for mobilization. Symbols such as blue birds, winter deer, and whale island—often punning on protest locations—circulated widely through reposts, providing a portable shorthand that reduced the costs of participation and strengthened collective identity (Shifman 2014). Yet

these same techniques were redeployed for attack: goblins, toads, and other AI-generated figures became tools of ridicule aimed at rival parties.

The dual use of symbolic AI imagery, mobilizing solidarity on the one hand and cloaking toxicity in cuteness on the other, demonstrates both the promise and the peril of generative aesthetics in contentious politics. This echoes scholarship on kawaii as a political aesthetic (Allison 2006; Ngai 2012) and extends recent work on meme culture and digital symbolism in protest movements (Shifman 2014; Milner 2016). We identify a new paradigm for out-group attacks called **kawaii toxicity**, which operates by packaging extreme toxicity in cute animals.

These findings carry broader implications for how digital protest repertoires evolve. Rather than displacing institutional channels, Threads communication was tightly coupled to them. The textual and visual strategies that spread most widely were those that connected online affect to offline mechanisms of accountability, particularly the unprecedented wave of recall elections. The Bluebird Movement thus shows how a new platform can bootstrap existing constitutional remedies, and transform symbolic protest into formal oversight and legislative action. This dynamic illustrates the continued importance of institutional context in shaping the trajectories of digital mobilization (Earl and Kimport 2011; Theocharis et al. 2015).

Our study has a few limitations. First, our observational design cannot establish causality, which is generally difficult without randomized experiments. Our single-platform, single-movement focus limits generalizability, and we measure circulation rather than mobilization outcomes or recall turnout. Future research should also consider how AI may facilitate misinformation, especially with Taiwan receiving some of the most misinformation in the world (Chang, Wang, and Fang 2024) and the growing literature on synthetic realities and fake news (Peng et al. 2025).

Acknowledgments

This research was supported by the Hanlon Scholars Program and Burke Research Initiative. We thank Hsuan Lo and Austin Wang for feedback.

References

Allison, A. 2006. *Millennial Monsters: Japanese Toys and the Global Imagination*. University of California Press.

AsiaPac. 2025. The new key to brand traffic: Threads ads will redefine social media with ‘the power of words’. Accessed September 10, 2025.

Bandy, J.; and Diakopoulos, N. 2021. More Accounts, Fewer Links: How Algorithmic Curation Impacts Media Exposure in Twitter Timelines. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–28.

Bennett, W. L.; and Segerberg, A. 2012. The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics. *Information, Communication & Society*, 15(5): 739–768.

Bonilla, Y.; and Rosa, J. 2015. #Ferguson: Digital Protest, Hashtag Ethnography, and the Racial Politics of Social Media in the United States. *American Ethnologist*, 42(1): 4–17.

Bozdag, E. 2013. Bias in Algorithmic Filtering and Personalization. *Ethics and Information Technology*, 15(3): 209–227.

Brady, W. J.; Wills, J. A.; Jost, J. T.; Tucker, J. A.; and Van Bavel, J. J. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28): 7313–7318.

Brandon, J. 2025. Meta Says Threads Now Has 400 Million Users, A Big Increase From April. *Forbes*. Accessed March 21, 2026.

Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 1877–1901.

Chang, H.; Shaman, B.; Chen, Y.-C.; Zha, M.; Noh, S.; Wei, C.; Weener, T.; and Magee, M. 2026. Generative Memesis: AI Mediates Political Information in the 2024 United States Presidential Election. *ICWSM’26 – 20th International AAAI Conference on Web and Social Media*.

Chang, H.-C. H.; and Fang, Y.-S. 2024. The 2024 Taiwanese presidential election on social media: Identity, policy, and affective virality. *PNAS Nexus*, 3(4): pgae130.

Chang, H.-C. H.; Haider, S.; and Ferrara, E. 2021. Digital civic participation and misinformation during the 2020 Taiwanese presidential election. *Media and Communication*, 9(1): 144–157.

Chang, H.-C. H.; Richardson, A.; and Ferrara, E. 2022. #justiceforgeorgefloyd: how instagram facilitated the 2020 black lives matter protests. *PLoS one*, 17(12): e0277864.

Chang, H.-C. H.; Wang, A. H.-E.; and Fang, Y. S. 2024. Us-skepticism and transnational conspiracy in the 2024 taiwanese presidential election. *Harvard Kennedy School Misinformation Review*.

Chen, T.; and Guestrin, C. 2016. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 785–794.

Dorogush, A. V.; Ershov, V.; and Gulin, A. 2018. CatBoost: gradient boosting with categorical features support. In *Proceedings of the Workshop on ML Systems at NeurIPS*. ArXiv:1810.11363.

Earl, J.; and Kimport, K. 2011. *Digitally Enabled Social Change: Activism in the Internet Age*. MIT Press.

Everington, K. 2024. Bluebird Movement raises NT\$2.6 million to post ‘Stand with Taiwan’ billboard in Times Square. Accessed September 9, 2025.

Freelon, D.; Marwick, A.; and Kreiss, D. 2020. False equivalencies: Online activism from left to right. *Science*, 369(6508): 1197–1201.

Friedman, J. H. 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5): 1189–1232.

- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. DeViSE: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, 2121–2129.
- Gao, T.; Fisch, A.; and Chen, D. 2021. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, 3816–3830.
- Gerbaudo, P. 2012. *Tweets and the Streets*. Pluto Press.
- Giles, D. C. 2002. Parasocial interaction: A review of the literature and a model for future research. *Media Psychology*, 4(3): 279–305.
- Gillespie, T. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.
- Hart, S. 1987. Shapley value. In Eatwell, J.; Milgate, M.; and Newman, P., eds., *Game Theory*, 210–216. Palgrave Macmillan.
- Hastie, T.; Tibshirani, R.; and Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition.
- Horton, D.; and Wohl, R. R. 1956. Mass communication and para-social interaction: Observations on intimacy at a distance. *Psychiatry*, 19(3): 215–229.
- Instagram. 2024. Continuing our approach to political content on Instagram and Threads. Official Instagram blog announcement.
- Jasper, J. M. 1998. The emotions of protest: Affective and reactive emotions in and around social movements. *Sociological Forum*, 13(3): 397–424.
- Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; and Liu, T.-Y. 2017. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 3146–3154.
- King, G.; Pan, J.; and Roberts, M. E. 2013. How Censorship in China Allows Government Criticism but Silences Collective Expression. *American Political Science Review*, 107(2): 326–343.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 4765–4774.
- McAdam, D.; Tarrow, S.; and Tilly, C. 2001. *Dynamics of Contention*. Cambridge University Press.
- Meta. 2023. Introducing Threads: A new way to share with text. Accessed September 9, 2025.
- Miller, D. 2017. *The Aesthetics of Cute Culture*. Routledge.
- Milner, R. M. 2016. *The World Made Meme: Public Conversations and Participatory Media*. MIT Press.
- Munger, K.; and Phillips, J. 2022. Right-Wing YouTube: A Supply and Demand Perspective. *The International Journal of Press/Politics*, 27(1): 186–219.
- Nachman, L. 2025. *Contested Taiwan: Sovereignty, Social Movements, and Party Formation*. University of Washington Press.
- Ng, J. Q. 2021. Crab Culture: Resisting Censorship through Homophonic Memes on Chinese Social Media. *Journal of Communication*, 71(6): 961–988.
- Ngai, S. 2012. *Our Aesthetic Categories: Zany, Cute, Interesting*. Harvard University Press.
- Palatucci, M.; Pomerleau, D.; Hinton, G. E.; and Mitchell, T. M. 2009. Zero-shot learning with semantic output codes. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 22, 1410–1418.
- Peng, Q.; Lu, Y.; Peng, Y.; Qian, S.; Liu, X.; and Shen, C. 2025. Crafting synthetic realities: Examining visual realism and misinformation potential of photorealistic ai-generated images. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 1–12.
- Penney, J. 2017. Social media and symbolic action: Exploring participation in the 2010 Red Shirt protests in Thailand. *New Media & Society*, 19(6): 901–918.
- Perelli, A.; and Bradley, S. 2025. An internal Meta planning deck reveals its early strategy for Threads, code-named Project 92. Accessed September 9, 2025.
- Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A. V.; and Gulin, A. 2018. CatBoost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 6638–6648.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. OpenAI Technical Report. https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Roberts, M. E. 2018. *Censored: Distraction and Diversion Inside China's Great Firewall*. Princeton University Press.
- Shifman, L. 2014. The Cultural Logic of Photo-Based Meme Genres. *Journal of Visual Culture*, 13(3): 340–358.
- Socher, R.; Ganjoo, M.; Manning, C. D.; and Ng, A. Y. 2013. Zero-shot learning through cross-modal transfer. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, 935–943.
- Theocharis, Y.; Lowe, W.; Van Deth, J. W.; and García-Albacete, G. 2015. Using Twitter to mobilize protest action: online mobilization patterns and action repertoires in the Occupy Wall Street, Indignados, and Aganaktismenoi movements. *Information, Communication & Society*, 18(2): 202–220.
- Tufekci, Z. 2017. *Twitter and Tear Gas*. Yale University Press.
- Tukachinsky, R.; and Stever, G. S. 2018. Theorizing development of parasocial engagement. *Communication Theory*.
- Wang, A. H.-E.; Wu, C. K. S.; Yeh, Y.-Y.; and Chen, F.-Y. 2024. Taiwan's protesters are against 'check-but-imbalance' and legislative overreach. Accessed September 9, 2025.

- Webster, J. G. 2014. *The Marketplace of Attention: How Audiences Take Shape in a Digital Age*. MIT Press.
- Weener, T.; Chang, H.; and Horiuchi, Y. 2025. Beyond cross-strait affairs: US-skeptical misinformation shapes policy preferences in Taiwan's 2024 election.
- Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-shot learning – The good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4582–4591.
- Yang, Z. 2024. Why Threads is Suddenly Popular in Taiwan. *MIT Technology Review*. Accessed: 2026-01-15.
- Yin, W.; Hay, J.; and Roth, D. 2019. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3914–3923.
- Zahn, M. 2023. 2 million sign-ups in 2 hours: What to know about newly launched Threads. Accessed September 9, 2025.
- Zha, M.; and Chang, H.-C. H. 2024. Gender Inequalities in Content Collaborations: Asymmetric Creator Synergy and Symmetric Audience Biases. arXiv:2411.05782.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? [Yes, please see the methods.](#)
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes, the abstract and introduction are scoped to reflect the results and contributions.](#)
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? [Yes, please see the methods.](#)
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? [Yes, please see the methods and the API’s sampling strategies.](#)
 - (e) Did you describe the limitations of your work? [Yes, please see the discussion.](#)
 - (f) Did you discuss any potential negative societal impacts of your work? Our research describes findings from a major social media platforms. They report on possible negative implications but the analysis itself is not directly applicable.
 - (g) Did you discuss any potential misuse of your work? Similar to above, there is no direct misuse given this is not an algorithm or specific technology. We make our interpretation clear in the discussion.
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? [Yes, please see the methods and discussion, and the API agreement with the platforms.](#)
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes.](#)
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? [Yes, please see the background and introduction.](#)
 - (b) Have you provided justifications for all theoretical results? [Yes, please see results and discussion.](#)
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? [Yes, the paper is set up as with dueling hypotheses.](#)
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? [Yes, please see the discussion where we also emphasize the need to consider other social media platforms.](#)
 - (e) Did you address potential biases or limitations in your theoretical framework? [Yes, please see the discussion.](#)
 - (f) Have you related your theoretical results to the existing literature in social science? [Yes, please see the literature review and discussion.](#)
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? [Yes, please see the discussion.](#)
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? [NA.](#)
 - (b) Did you include complete proofs of all theoretical results? [NA.](#)
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Code for replication will be included upon publication.](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Please see methods.](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [NA.](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [NA.](#)
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? [NA.](#)
 - (f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? [We manually validate the classification to minimize the possible error. Please see methods.](#)
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? [NA.](#)
 - (b) Did you mention the license of the assets? [NA.](#)
 - (c) Did you include any new assets in the supplemental material or as a URL? [NA.](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [Please see methods.](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Please see methods.](#)
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see ?)? [NA.](#)
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see ?)? [NA.](#)
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...

- (a) Did you include the full text of instructions given to participants and screenshots? *NA*.
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? *Yes, please see methods*.
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? *NA*.
- (d) Did you discuss how data is stored, shared, and de-identified? *Yes, please see methods*.

Prompt for visual annotation

The full prompt can be found in the [GitHub](#).

```

1 As a political respondent analyzing the
  domestic politics of Taiwan and
  international politics, your goal is
  to assess and provide accurate
  descriptions of images.
2
3 ### Instructions
4
5 The following formatting notes are
  extremely important to follow exactly
  correctly:
6 - If an image includes any text, please
  extract it all into a string.
7 - Please give probabilities as
  percentage likelihood (i.e. 1% if
  very unlikely and 99% if extremely
  likely).
8 - Corresponding to the questions below,
  you will need to output a JSON object
  .
9 - Return the structured JSON only, with
  no additional text, descriptions, or
  explanations.
10
11 ### Questions
12
13 This is an image that I want to upload.
  Describe what is in it.
14
15 === Basic Identification ===
16 - Classify the image into one of the
  following typologies:
17 (Map, Logo, Infographic, Weather Map,
  News Footage, Legislative Footage,
  Merchandise, Art, Cartoon, Personal
  , Nature). If none fit, use "other
  ". (label as Category:)
18 - Create a description of the image,
  limited to 60 words. (label as
  Description:)
19 - Classify the image into one of the
  following visual content categories:
20 (Crowd-Based Protest Photos, Non-Crowd
  -Protest Human Photos, Non-Human
  Photos, Non-Photograph Visuals).
21
22 === Textual Content ===
23 - If the image includes any text,
  extract it into a string. (label as
  Text:)
24 - If it contains text, what language is
  it? If multiple languages, list all.
  (label as Language:)
25 - If the image includes any hashtags,
  extract them. (label as Hashtags:)
26 - If the image includes any locations,
  extract them. (label as Location:)
27
28 === Visual Content ===
29 - What is the probability the image
  contains a blue bird? (label as
  BlueBird:)
```

```

30 - What is the probability the image
    contains animals other than blue
    birds? (label as Animal:)
31 - If animals are present, what kind? (
    label as AnimalType:)
32
33 === In-Group/Out-Group Partisan
    Strategies ===
34 - What is the probability the image
    supports the DPP? (label as ProDPP:)
35 - Attacks the DPP? (label as AgainstDPP
    :)
36 - Supports the KMT? (label as ProKMT:)
37 - Attacks the KMT? (label as AgainstKMT
    :)
38
39 {
40   "Category": string, // One of: Map,
    Logo, Infographic, Weather Map,
    News Footage, Legislative Footage,
    Merchandise, Art, Cartoon, Personal
    , Nature, other
41   "Description": string, // ≤ 20-word
    description
42   "VisualCategory": string, // One of:
    Crowd-Based Protest Photos, Non-
    Crowd-Protest Human Photos, Non-
    Human Photos, Non-Photograph
    Visuals, other
43
44   "Text": string, // Full extracted text
    from image, if any
45   "Language": [string], // List of
    languages (e.g., ["zh", "en"])
46   "Hashtags": string,
47   "Location": string,
48
49   "BlueBird": int, // 0-100 (%)
50   "BlueBirdFly": int,
51   "Animal": int,
52   "AnimalType": string,
53
54   "ProDPP": int,
55   "AgainstDPP": int,
56   "ProKMT": int,
57   "AgainstKMT": int,
58
59 }

```

Accuracy for stratified sampled labeling

Full accuracy scores can be found in the Open Science Framework files for this project: https://osf.io/j5e3q/overview?view_only=11ee872ae1ff4b90801ee38809dcc356.

Feature	Value
VisualCategory	91.9%
Photograph	95.2%
Taiwan	78.6%
Cute	100%
InternationalPolitics	92.1%
DomesticPolitics	90.5%
Meme	100%

Table 2: Model feature performance summary.

Checks for Colinearity

Variable	VIF
Intercept	5.18
Photograph	1.06
Taiwan	1.88
Cute	1.22
International Politics	1.33
Domestic Politics	2.23
Meme	1.18
Human Numbers	1.00
DPP	1.23
KMT	1.04
TPP	1.02

Table 3: Variance Inflation Factors (VIF) for Visual-Based Predictors.

Variable	VIF
Intercept	12.00
Formality	1.23
Sentiment	1.30
Domestic Politics (Text)	1.51
Political Rally (Text)	1.14
Against KMT	1.33
Recall (Text)	1.11
Pro-DPP	1.16
Against DPP	1.30
Pro-KMT	1.07
Pro-TPP	1.05
Against TPP	1.15

Table 4: Variance Inflation Factors (VIF) for Text-Based Predictors.