國立臺灣大學管理學院資訊管理學研究所

碩士論文

Department of Information Management
College of Management
National Taiwan University

Master Thesis

社群媒體平台合作影片之競食效應預測

Is Your Guest an Ally or an Enemy? Predicting Cannibalization Effects of Featured Videos on Social Media Platforms

林璟耀

Ching-Yao Lin

指導教授：魏志平 博士

Advisor: Chih-Ping Wei, Ph.D.

中華民國 112 年 7 月

July 2023

doi:10.6342/NTU202304164

誌謝

在這段碩士生涯旅程即將結束之際，我想要衷心地表達我最深的謝意和感激之情。在我完成這篇論文的過程中，有著無數的陪伴和支持，讓我能夠克服種種困難，順利地走到這一步。首先，我要對魏志平教授表示最衷心的感謝。您不僅給予我們無限的靈感和寶貴的建議、讓我們收穫良多，更在我們遇到疑惑和困難的時候，不辭辛勞地指導、耐心解答我們的疑問。那些半夜的會議，讓我們感受到您對我們的關心和支持，辛苦您了！您的教導和鼓勵，是我們前進的動力。

再來我要感謝我的同學們，你們是我這段時光中最重要的陪伴。宇雋、海因、佩琳、雨柔、佑甄，我們彼此鼓勵、共同進步、互相協助，也有無數歡笑的時候。尤其感謝宇雋，在我論文的最後階段和我在實驗室一同度過的時光，你給我的各種幫忙對我來說意義非凡。我也要感謝實驗室學長姐，思涵、沛原、仁瀚，尤其感謝仁瀚學長跟我開了好多次會，提供我各種想法跟技術上的幫助。

我的家人，您們是我最堅強的後盾。感謝您們無私的栽培，讓我可以放心自由地追求更高的目標。謝謝您們一直以來的支持和鼓勵，是您們的愛讓我能夠堅持不懈，走到今天這個重要的里程碑。您們的支持是我無限的動力。此外，我也要感謝其他好朋友們。文楷、祐琳、書戎、秉諮，感謝你們陪伴我度過了最後的時光。你們的陪伴和鼓勵，讓我在研究的路上不再孤單。我們一同面對挑戰、一同歡笑、一同前進。

最後，感謝所有在背後默默支持我的人們。這個論文的完成，離不開你們的幫助和鼓勵。這個論文不僅是我的成果，更是我們共同的努力和付出。在論文的撰寫過程中，我學到了無數的知識和經驗，也經歷了無數的挑戰和成長。這段時光將成為我人生中寶貴的回憶。謝謝大家，讓我有機會完成這個重要的任務，我會繼續保持學習和進步，並且不會忘記您們的愛。

林璟耀 謹誌

于國立臺灣大學資訊管理研究所

中華民國一一二年八月

doi:10.6342/NTU202304164

摘要

傳統上，關於競食效應(Cannibalization)的討論通常局限於產品的範疇。 在本研究中，我們提出了一個新穎的研究任務，探討社群媒體網紅的合作影片 中的競食效應預測。 本研究旨在預測社群媒體網紅邀請其他網紅嘉賓出現自己的影片中時，導致其自身未來影片的平均觀看量下降的情況，稱之為競食效應。

為了處理這個研究任務，我們提出一種新穎的深度神經網路預測模型 Cannibalization Identification with Influencer Encoder (CIIE)，其利用網紅編碼器(Influencer Encoder)萃取主持網紅、特邀嘉賓網紅，以及他們過去創作內容的關鍵資訊。此外，我們也提出多種基準模型(Baseline Models)以綜合評估我們提出的 CIIE 模型之整體表現，其中，包括先驗機率模型 (PPM)、有約束的先驗機率模型(CPPM)和隨機預測模型(RGM)。根據我們的實驗結果，我們提出的 CIIE 模型之在所有方法中表現最優，尤其對於預測少數類別方面表現尤其出眾，這也是我們研究的主要關注點。這項研究為我們對社群媒體網紅競食效應的理解做出了具體的貢獻，並證實我們提出的深度神經網路預測模型可以有效地預測可能發生競食效應之社群媒體網紅的合作影片。

關鍵字: 深度學習、競食效應、合作影片、社群媒體、網紅、網紅編碼器

doi:10.6342/NTU202304164

Abstract

Traditional discussion regarding cannibalization is restricted to the context of products. In this research, we present a novel research task which investigates the prediction of the cannibalization effect in the context of featured videos by social media influencers. This research aims to identify instances where the exposure of a social media influencer as a guest in another influencer's video leads to a decline in their own video's viewership, known as cannibalization.

To address this research task, a novel deep neural network predictive model, referred to as Cannibalization Identification with Influencer Encoders (CIIE), is proposed, utilizing influencer encoders to capture essential information about both the host and guest influencers and their past video content. The model's effectiveness is evaluated against various benchmark methods, including Prior Probabilistic Model (PPM), Constrained Prior Probabilistic Model (CPPM), and Random Guess Model (RGM). According to our evaluation results, our proposed CIIE model outperforms all benchmarks and is especially effective in the predict minority classes, which is the main focus of our study. This research contributes to a comprehensive understanding of cannibalization among social media influencers and underscores the potential of our proposed DNN model as a valuable tool for predicting possible cannibalization effects for featured videos.

Keywords: Deep learning, Cannibalization, Featured Videos, Social Media, Influencers, Influencer Encoders

doi:10.6342/NTU202304164

Table of Contents

doi:10.6342/NTU202304164

doi:10.6342/NTU202304164

doi:10.6342/NTU202304164

List of Tables

doi:10.6342/NTU202304164

List of Figures

doi:10.6342/NTU202304164

Chapter 1 Introduction

1.1 Background

Social media applications, commonly referred to as social apps, have assumed unprecedented importance in people's lives. The demand for these platforms is staggering, with approximately 60% of the world's population actively using social media, dedicating an average of 2 hours and 24 minutes each day on these applications[1]. Deloitte's report further reveals that over 50% of U.S. consumers access their social media apps on a daily basis[2]. Evidently, social media has become an integral part of daily routines for a vast majority of individuals.

Notably, social media platforms are not only indispensable for personal use but also wield immense influence in the business landscape. A remarkable 72.8% of internet users employ social media for brand research[1]. This shift in consumer behavior has prompted businesses to recognize the potential of these platforms as powerful tools for brand awareness and customer engagement. As a result, a substantial 81% of organizations leverage social media to raise brand awareness, and 71% of B2C organizations have reported to actively adopt influencer marketing strategies[1]. At the core of these strategies are social media influencers, who stand as independent third-party spokespersons with the ability to shape and alter audience perspectives through blogs, tweets, and various other social media outlets (Freberg et al., 2011).

---

[1] Kemp, S. (2023). Digital 2023 April Global Statshot Report (April 27, 2019). Available at https://datareportal.com/reports/digital-2023-april-global-statshot (Retrieved on July 6, 2023).

[2] Deloitte. (2016). Digital democracy survey: A multi-generational view of consumer technology, media and telecom trends (tenth edition). Deloitte. Available at https://www2.deloitte.com/content/dam/Deloitte/za/Documents/technology-media-telecommunications/ZA_Deloitte_Digital_Democracy_Survey_Final.pdf (Received on July 6, 2023).

doi:10.6342/NTU202304164

The exponential growth of influencer marketing is a testament to the increasing reliance on social media for business communication. As of 2022, the global influencer marketing market value reached a staggering 16.4 billion U.S. dollars, more than doubling since 2019[3]. Indeed, social media platforms have proven to be effective tools in achieving marketing objectives, particularly in areas such as customer relationship management and customer involvement (Alalwan et al., 2016). It comes as no surprise to witness that the burgeoning popularity of influencer marketing has rendered platforms such as YouTube exceptionally suitable channels for the implementation of such marketing strategies (Xiao, Wang, & Chan-Olmsted, 2018). Such substantial investment underscores the pivotal role that social media platforms play in facilitating interactions between businesses and their customers.

Among the plethora of social media applications, video streaming platforms stand out as a significant player. Deloitte's report revealed that nearly one out of every two U.S. consumers subscribes to at least one video streaming service[2]. Furthermore, a staggering 91.8% of internet users regularly watch online video content on a weekly basis[1]. Video-based platforms, such as TikTok and YouTube, garner notably higher user engagement compared to other social apps, solidifying their positions as the top two platforms in terms of time spent on social apps[1].

The impact of social media video streaming platforms extends beyond individual users; businesses and organizations also leverage these platforms to their advantage. A recent study has shown that 80.5% B2B decision-makers use video streaming sites to research work-related purchases, and 16.3% B2B decision makers say video streaming

[3] Dencheva, V. (2023). Available at Influencer marketing worldwide - statistics & facts (January 19, 2023). Statista. https://www.statista.com/topics/2496/influence-marketing (Received on June 30, 2023)

doi:10.6342/NTU202304164

sites influence their ultimate purchase decisions[1]. These highlights the role of video content in influencing B2B buying decisions. Moreover, the efficacy of social media video streaming platforms for marketing purposes is also undeniable. A considerable 75% of marketers use YouTube influencers as a part of their marketing strategy; the average ROI is estimated around $6.50 earned for every $1 spent[4]. Certainly, social media video streaming platforms are substantial in modern marketing effort and are a key driver in business nowadays.

Influencer marketing has emerged as a highly lucrative business on social media video streaming platforms, such as YouTube and TikTok, due to its demonstrated effectiveness. Influencers in these platforms have the persuasive power to influence and shape consumers' attitude through their uploaded videos. Therefore, it comes as no surprise to see that more and more brands seek to utilize this mechanism to disseminate their marketing or advertising messages. For example, the YouTube influencer market size was valued at around $6 billion in 2020[1].

Given the significant market potential, influencers on these platforms derive a substantial portion of their income from advertising and sponsorship revenues, as they assist brands in connecting with potential customers. For instance, according to BBC, YouTubers charge brands an average of $187,500 per sponsored video[5]. Similarly, TikTokers earn between $100,000 and $250,000 for a branded video, as reported by Forbes in 2022[6].

---

[4] Gitnux (2023). The Most Surprising YouTube Influencer Marketing Statistics and Trends in 2023 (April 5, 2023). Gitnux. Available at https://blog.gitnux.com/youtube-influencer-marketing-statistics/ (Received on July 2, 2023)
[5] BBC (2016). Celebrity YouTube promotion fee '$187,000 on average' (August 31, 2016). Available at https://www.bbc.com/news/technology-37234385 (Received on July 2, 2023)
[6] Brown, A. & Freeman, A. (2022). Top-Earning TikTok-ers 2022: Charli And Dixie D'Amelio And Addison Rae Expand Fame—And Paydays. Forbes. Available at https://www.forbes.com/sites/abrambrown/2022/01/07/top-earning-tiktokers-charli-dixie-damelio-addison-rae-bella-poarch-josh-richards/?sh=ef32f123afa4 (Received on July 4, 2023)

doi:10.6342/NTU202304164

In pursuit of maximizing their income, social media influencers on video streaming platforms strive to enhance their popularity. The correlation between an influencer's popularity and a business's ability to reach and impact a larger audience through their videos directly translates to increased sales and overall revenue for brands. Consequently, influencers are paid based on their popularity, which is often measured by the total number of subscribers or followers, or the number of views of their recent videos. Therefore, influencers in these platforms adopt a variety of strategies and tactics to heighten their popularity.

1.2 Research Motivation

Among the various methods influencers employ to enhance their popularity, collaboration with other influencers stands out as a prominent strategy on social video streaming platforms. This collaborative approach is instrumental in increasing an influencer's earnings, as it allows them to expose their videos to a broader audience, ultimately leading to an upsurge in video views (Koch et al., 2018).

The essence of this collaboration lies in one influencer, referred to as the host (Influencer A), featuring another influencer, the guest (Influencer B), in a video on the host's channel. The title of this collaborated video includes a phrase like "feat. B", "ft. B", or "@B" to indicate their collaboration and this collaborated video is called a featured video. It is hosted on A's channel and therefore creates an opportunity for A's viewers to be introduced to B. The rationale is that the viewers of A are more likely to watch B's videos and even become subscribers after the introduction of B in A's featured video, creating a cross-pollination effect. This reciprocal activity is tested to be effective in

doi:10.6342/NTU202304164

general in fostering mutual growth and expansion of their respective channels (Koch et al., 2018; Ma, Gui, & Kou, 2023).

However, such collaborations are not always beneficial to both sides; there may be a significant downside, despite the potential benefits of collaborations between influencers, that can be detrimental to social media businesses—cannibalization. It can cause unintended harm to a social media business when pursuing such a collaboration strategy. Cannibalization, in the context of social media influencer marketing, refers to the phenomenon wherein the introduction of a new channel or collaboration might lead to the redistribution of the existing audience across various channels. This audience fragmentation could result in a decline in viewership and engagement for the original influencer, thereby diminishing their overall effectiveness as a brand ambassador or spokesperson.

Building upon the previous example, it is possible that followers of influencer A are drawn to another influencer B to such an extent that it significantly diminishes A's popularity, often observed by a nosedive in total views or watch time. This undoubtedly goes against the original intention of collaborating with B and hampers A's effectiveness as a brand ambassador or spokesperson. Naturally, influencers make concerted efforts to avoid these outcomes. To our surprise, however, in the dataset that we collect, around 3.3% of featured videos are considered as cannibalized. This finding highlights the presence and significance of cannibalization as a genuine concern in the realm of influencer marketing.

Undoubtedly, cannibalization poses a significant risk that necessitates proactive identification and prevention measures. The ability to predict potential cannibalization before making decisions about featuring another channel is of paramount importance. By

doi:10.6342/NTU202304164

doing so, we can avert detrimental collaborations that might otherwise compromise the popularity and intrinsic value of an influencer's channel. Hence, the anticipation and avoidance of cannibalization are vital for sustaining the long-term success and relevance of social media influencers.

However, cannibalization prediction remains an unexplored area in extant research. The identification of cannibalization is often insufficiently addressed, with discussions mainly confined to product contexts, focusing on the diversion of sales between existing and new products within a single company. To the best of our knowledge, the cannibalistic effect resulting from collaborations among social media influencers has not been investigated to date. As a consequence, no prior attempts have been made to predict cannibalization, especially concerning featured videos on social media video streaming platforms. There is a dearth of knowledge and established methodologies to foresee and address the risks associated with cannibalization in such contexts.

Overall, we believe that further investigation and research in this domain are essential to comprehensively grasp the dynamics of cannibalization and its implications on influencer marketing strategies. By doing so, influencer marketing campaigns can be strategically planned to promote positive outcomes for all parties involved and mitigate the risks associated with diverted viewership and declining channel popularity. This pursuit of knowledge will undoubtedly empower marketers and influencers alike to make informed decisions, ensuring sustainable and successful engagements that resonate with audiences and drive value for brands and content creators. As the landscape of social media and influencer marketing continues to evolve, staying vigilant in our exploration of cannibalization can pave the way for more effective, efficient, and mutually beneficial collaborations in the future.

doi:10.6342/NTU202304164

1.3 Research Objectives

One primary objective of this thesis is to define cannibalization in the context of featured videos within social apps and establish a prediction task for identifying this phenomenon in advance. Cannibalization, in this study, refers to instances where the popularity of the host channel experiences a significant drop following collaboration with an invited guest influencer. By defining and predicting cannibalization, this research aims to shed light on unintended outcomes that may occur as a result of influencer collaborations. Our task is defined as a three-class classification problem, where each featured video is classified into one of the three classes: cannibalized, boosted, and unaffected. The "cannibalized" class represents featured videos where their hosts' viewership declines significantly after collaboration, while the "boosted" class signifies featured videos that result in a substantial increase in the hosts' viewership. The "unaffected" class includes featured videos with no significant change in viewership observed.

Another primary objective of this thesis is to develop a predictive model that can identify whether a featured video will cannibalize the host's viewership before the collaboration commences. Leveraging the current state of channels and videos from both the host and guest influencers, this research seeks to develop an effective cannibalization prediction method. We will develop an influencer encoder to learn the embeddings for each influencer, including both the host and the guest(s), and utilize these embeddings to make predictions. This model will enable informed decision-making for host influencers

doi:10.6342/NTU202304164

contemplating collaborations with guest influencers, thereby optimizing the outcome of their collaborations.

To address the prominent problem of cannibalization prediction, this research proposes a novel DNN model, namely, Cannibalization Identification with Influencer Encoders (CIIE), that can effectively capture the intricate patterns and dynamics underlying cannibalization occurrences in social media video streaming platforms. To facilitate the investigation of cannibalization prediction, this research endeavors to collect a dataset of featured videos and relevant data. This dataset will encompass a representative sample of social app channels and influencer collaborations, enabling thorough analysis and model development. To validate the effectiveness of the proposed CIIE model, we aim to establish three types of benchmarks for evaluation. Utilizing stratified cross-validation, the research will assess the model's performance against these benchmarks, providing insights into the model's effectiveness and potential real-world applicability.

In conclusion, this work seeks to make novel contributions to the understanding of cannibalization in the context of featured videos in social media video streaming platforms. The outcomes of this research are intended to assist influencers in making informed collaboration decisions and optimizing the success of their partnerships in the evolving landscape of influencer marketing.

doi:10.6342/NTU202304164

Chapter 2 Literature Review

2.1 Traditional Definition of Cannibalization

Cannibalization is a concept that lacks a universally agreed-upon definition (Lomax et al., 1996), and existing definitions have primarily centered around cannibalization among products within a company. An early definition by Heskett (1976) described cannibalization as “the process by which a new product gains sales by diverting sales from an existing product.” This merely depicted the general concept but not highlighting the essential distinction that cannibalization involves products from the same company. Copulsky (1976) addressed this limitation by defining cannibalization as “the extent to which one product’s sales are at the expense of other products offered by the same firm.” Another definition by Kerin, Harvey and Rothe (1978) introduced a slightly different perspective, “‘Redistributed’ revenue, in that existing buyers are substituting one item for another in the company’s product portfolio.” This definition shifted the focus to revenue rather than sales.

Novelli (2013) offered a more detailed definition of cannibalization after summarizing most frequently cited definitions. According to his definition, cannibalization is an intra-firm sales diversion phenomenon, where a particular product or service (the cannibal) draws in sales by redirecting potential sales that would have otherwise gone to another product or service (the victim) within the same organization that collects both of their revenues. This definition is fairly detailed and provides valuable insights into the concept. Yet, it is noteworthy that this definition, like many others, remains confined to the context of products within a company.

doi:10.6342/NTU202304164

As we delve deeper into the existing definitions of cannibalization, it becomes evident that they predominantly revolve around product-level cannibalization within a single company. While these definitions provide valuable insights into the concept's traditional context, they fail to encompass the intricacies that arise when cannibalization occurs during brand alliances, where brands come from different companies. This limitation becomes particularly crucial when considering the rising prevalence of influencer marketing and collaborations on social media video streaming platforms.

Another notable aspect of the current state of cannibalization definitions is the lack of clear characterizations to distinguish between different cannibalization types and the absence of standardized measurement techniques to accurately gauge their impact. This is understandable given the complexity of cannibalization, and as a result, researchers often have the autonomy to prioritize specific aspects based on their investigation, such as sales volume or revenue in the context of product cannibalization. However, marketers and researchers may face challenges in determining the extent of cannibalization and its implications for their marketing strategies and brand positioning across various contexts.

Therefore, to address these limitations and to clear the ambiguity in current definitions, three essential steps are needed. First, it is imperative to draw a clear contrast between cannibalization among products within a company and that among brands during co-branding efforts. This differentiation will help researchers and marketers better grasp the nuances and consequences associated with each type of cannibalization, enabling them to tailor strategies accordingly. Second, identification of standard measurements for distinct types of cannibalizations is of significance. By establishing standardized metrics and criteria for measurement, researchers can more effectively evaluate the impact of cannibalization, regardless of its type, leading to clearer insights and a better

doi:10.6342/NTU202304164

understanding of its implications in various contexts. Finally, characterizing each type of cannibalization with clear-cut standards is crucial for developing a comprehensive understanding of the phenomenon.

The first essential step, as mentioned, is to recognize that cannibalization can be categorized into two main types based on its underlying causes. The first type is *product cannibalization*, which occurs when different products from the same company compete with each other. In this scenario, one product captures a portion of the market share or revenue stream that would have otherwise gone to another existing product of the same company. The second type is *co-branding cannibalization*, which is caused by brand alliances and typically involves two or more companies. A brand alliance, also known as co-branding, is defined as two or more existing brands are combined into a joint product or are marketed together in some fashion (Keller, 2013; Rao, Qu, & Ruekert, 1999; Ruekert & Rao, 1994; Simonin & Ruth, 1998). In this case, one brand may cannibalize another brand involved in the co-branding initiative, where both brands are placed within the same marketing context, such as an advertisement, a promotion, or a joint product.

The second step is to acknowledge the different measurement for different types of cannibalizations. For product cannibalization, the most common metrics are sales volume and sales value. Researchers typically analyze changes in absolute value or (relative) market share to assess the impact of one product on the sales of another within the same company. On the other hand, co-branding cannibalization can be evaluated using two distinct methods. The first method is to examine the decrease in the monetary value of a brand, which is typically evaluated through revenue changes. This measurement is akin to the concept of sales-based brand equity (SBBE), defined as "the revenue premiums generated by a brand compared to an essentially identical but unbranded offering"

doi:10.6342/NTU202304164

(Koschmann, 2017). The second method is to detect a denigration of customers' perception toward a brand (typically assessed via customer surveys), also known as consumer-based brand equity (CBBE) in abundant prior studies (Keller, 1993, 2016; Koschmann, 2017; Lassar, Mittal, & Sharma, 1995; Shocker, Srivastava, & Ruekert, 1994).

The last step is to characterize each type of cannibalization with clear-cut standards, which involve three key components, as identified by Novelli (2013). These are organized by different types of cannibalizations in Table 1. While these components originally focused on delineating product cannibalization, we believe these components can be appropriately adapted and extended to encompass co-branding cannibalization as well.

For product cannibalization, the first component, the economic entities participating in the process, refers to the cannibalizing product (the cannibal) and the cannibalized existing product (the victim). These are the products within the same company that compete with each other for market share and revenue. The second component relates to a common resource base, which in this case refers to the sales volume or income stream of the same company. Finally, the third component is a specific relationship linking the value-generating process. In the context of product cannibalization, the relationship is a shared marketplace, where both the cannibalizing and cannibalized products are presented to the same group or overlapped groups of customers, allowing them to choose between the two products.

For co-branding cannibalization, the corresponding three components are the two brands forming a brand alliance, changes in brand equity, which can be evaluated through either SBBE or CBBE, and the channel through which customers are exposed to the brand

doi:10.6342/NTU202304164

alliance, such as a multi-brand product or a co-branding advertisement that potentially alters customers' purchase decisions.

Table 1: Three Components that Characterize Cannibalization

	Economic Entities	**Common Resource**	**Specific Relationship Linking the Value-generating Process**
Product Cannibalization	Products	Sales volume or revenue	A shared marketplace that displays both products to the same batch of customers
Co-branding Cannibalization	Brands	CBBE or SBBE	The channel whereby customers are exposed to their alliance (e.g., multi-brand products or co-branding advertisements)
Influencer Cannibalization	Influencers	Video views or subscriber count	The social platform where viewers can access the influencers' featured videos

2.2 Cannibalization Among Social Media Influencers

When considering cannibalization in the context of featured videos in social media video streaming platforms, it bears a closer resemblance to co-branding cannibalization. In this scenario, social media influencers can be viewed as distinct brands, each possessing their unique identity and followers, while the videos they produce become their products, which are offered to their respective audiences. A collaboration between these influencers corresponds to a brand alliance, where a featured video becomes a joint product that combines their individual influence and content.

The essence of cannibalization in this context lies in influencers potentially stealing each other's popularity. In other words, the introduction of a featured video could result in one influencer gaining popularity at the expense of the other. The measurement of this cannibalization effect is often tied to metrics such as the number of subscribers or total video views, which serve as vital indicators of an influencer's reach and influence. These

doi:10.6342/NTU202304164

metrics act as proxies for the preferences and perceptions of their followers, which is similar to CBBE mentioned above.

For cannibalization among social media influencers, which we refer to as influencer cannibalization, the three key components can be summarized as follows, as presented in Table 1. The first component is the collaborating influencers involved in the featured video—the host and the guest(s)—whose unique individualities and content contributions combine to create a co-branding-like experience for their audiences. The second component is changes in popularity of the involved influencers, which can be measured through metrics such as total video views or subscriber count. These popularity indicators offer valuable insights into the level of engagement and interest among their respective audiences. Lastly, the specific relationship linking the value-generating process (the third component) is the social media platform where viewers are exposed to the influencers' featured videos.

Additionally, it is important to note that, in the realm of co-branding, the concept of brand fit holds significant importance, as it directly contributes to the overall success of collaborative ventures. Many prior studies have shown that relationship between brands has a considerable impact on the success of co-branding (Newmeyer, Venkatesh, & Chatterjee, 2014; Paydas Turan, 2021), and co-brand similarities impact brand fit perception (Ahn, Kim, & Sung, 2021; Bouten, Snelders, & Hultink, 2011; Decker & Baade, 2016; Simonin & Ruth, 1998; Van der Lans, Van den Bergh, & Dieleman, 2014).

Therefore, acknowledging the significance of brand compatibility, we have integrated this factor into our analysis. To achieve this, we employ a method that involves leveraging the profiles of the preceding *k* videos of the host influencer and the guest influencer(s). These profiles serve as representative indicators of the brand fit and

doi:10.6342/NTU202304164

compatibility between the involved entities. This approach enables us to effectively capture and quantify the level of alignment between brands in a co-branding context.

As such, putting everything together, our operational definition of cannibalization in the context of featured videos on social media video streaming platforms is defined as "a significant percentage drop in the host channel's total video views of *k* videos after a featured video, as compared to that of *k* videos before that featured video." In our research, we have chosen to utilize total video views as the primary measurement of cannibalization, as opposed to subscriber count, due to two compelling reasons.

First, subscriber count is not considered a practical measurement for assessing cannibalization since it tends to remain relatively stable or may even continue to increase over time. This is because followers who are no longer interested in an influencer can simply stop clicking into his/her videos, leading to reduced engagement without necessarily affecting the total number of subscribers. Consequently, social media platforms' recommendation algorithms cease showing relevant videos to disengaged followers, further complicating the assessment of cannibalization using subscriber count.

On the other hand, total video views offer a more dynamic and real-time reflection of an influencer's popularity. Viewing a video is analogous to consuming a product, and as the consumption of videos increases, it directly correlates to an influencer's rising popularity. By focusing on total video views, we can capture the actual engagement and interest of the audience, providing a more accurate assessment of the impact of cannibalization on an influencer's content and overall appeal.

Therefore, through our chosen measurement of total video views, we aim to better capture the dynamics of cannibalization within the realm of featured videos on social

doi:10.6342/NTU202304164

media platforms, thereby contributing to a more accurate analysis of influencer collaborative strategies.

2.3 Summary

By reviewing the literature related to cannibalization, we discover that cannibalization lacks a universally agreed-upon definition, and existing ones have predominantly centered around product cannibalization within a company. However, it is essential to differentiate between different types of cannibalizations, particularly when considering brand alliances and collaborations involving influencers on social media platforms. While traditional definitions offer valuable insights, they may not fully encompass the complexities arising from such modern marketing practices.

From our literature review, we learn that influencer cannibalization closely resembles co-branding cannibalization. The essence of influencer cannibalization consists in the potential for one influencer to gain popularity at the expense of the other. This effect is often measured by metrics such as total video views or subscriber count; however, to better capture the dynamics of influencer cannibalization, our operational definition focuses on total video views rather than subscriber count.

doi:10.6342/NTU202304164

Chapter 3 Our Proposed CIIE Model for Cannibalization Prediction

3.1 Problem Formulation

Our task is a three-class classification problem, where we aim to predict the impact of a given featured video on the performance of the host's next k videos. The three possible outcomes are as follows. (1) Cannibalized: The featured video results in a significant percentage drop in views (θ %) for the host's next k videos. This indicates that the audience's attention has shifted away from the host's subsequent content significantly due to the impact of the featured video. (2) Boosted: The featured video leads to a substantial increase in views (-θ %) for the host's next k videos. This suggests that the featured video's success has positively influenced the host's subsequent content, attracting more viewers and engagement. (3) Unaffected: The featured video has no significant impact on the views of the host's next k videos, and the viewership remains relatively stable (between -θ % and θ %). This outcome indicates that the featured video and the host's subsequent content have independent viewership without affecting each other.

To make a prediction, we will utilize channel features from both the host and guest channels, as well as the features of the host's past k videos and the guest's past k videos leading up to the featured video. These data points will be leveraged to build a classification model that can determine whether the collaboration is cannibalized, is boosted, or has no significant effect on the host.

doi:10.6342/NTU202304164

To illustrate, consider a case where there is only one host and one guest channel, as depicted in Figure 1. In this figure, each vertical bar represents a video and the height of the bars represents their respective video views. Thus, the taller the bar, the higher the corresponding video views. Also, note that the upper section of the figure indicates the videos owned and published by the host, while the lower section is demonstrative of those owned and published by the guest. The two x-axes resting at the bottom of the bars respectively indicates the order of the host's and the guest's videos' published time; therefore, bars toward the left stands for videos published earlier, and those to the right later. Finally, the red bar is a featured video, where the host invites the guest to appear in it.

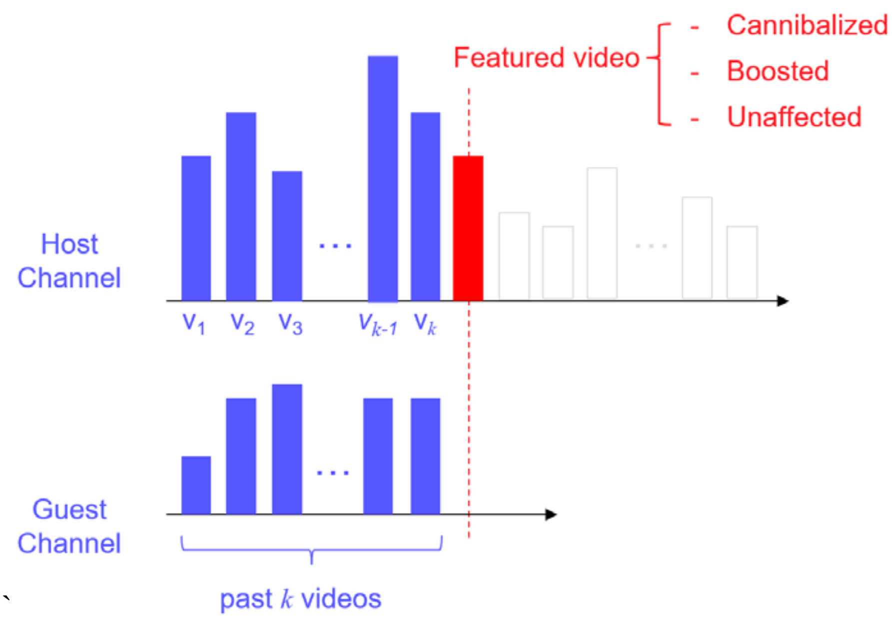

Figure 1: An Illustration of Our Cannibalization Prediction Task

As mentioned, this featured video will be classified into one of the following: cannibalized, boosted, or unaffected. This is done by feeding four types of inputs into our model: (1) the host's channel data, including its total views, subscriber count, channel

doi:10.6342/NTU202304164

title, etc, (2) the guest's channel data, (3) the host's past k video data, including views, likes, video title, etc, and (4) the guest's past k video's data. Such a scenario is to simulate the circumstance where the host is deciding whether to feature the guest in his or her next video. At this juncture, all one has is past information; the performance of the next k videos following the featured video (red bar) is not known.

3.2 Model Architecture

3.2.1 Overview

Suppose that we have a featured video labeled as one of the three classes, namely, cannibalized, boosted, or unaffected. Then, to make our prediction, we first have to gather channel features of the host and the guest channels of that featured video, and video features of the past k videos leading up to the featured video from both the host and all the guests. These are raw inputs to our proposed model architecture.

Then, for each influencer, these raw inputs are fed into his or her respective influencer encoder. The function of the influencer encoders, including a host and a guest influencer encoder, is to summarize all the information of an influencer and generate a representation for the focal influencer. Accordingly, we will obtain a host representation and potentially multiple guest representations.

Note that the guest influencer encoder is shared by all guests. This means that all guests are fed into the same influencer encoder to produce their individual representations. Also note that the number of guests can vary for different featured videos. Accordingly, we take the mean across all the produced guest representations to get a general guest representation that carries the overall information and impact of the guests.

doi:10.6342/NTU202304164

Finally, the host representation and the average guest representation are fed into a simple classifier to output a probability distribution of the three classes. The prediction is evaluated against the ground truth of the featured video. The model structure described in this section is illustrated in Figure 2.

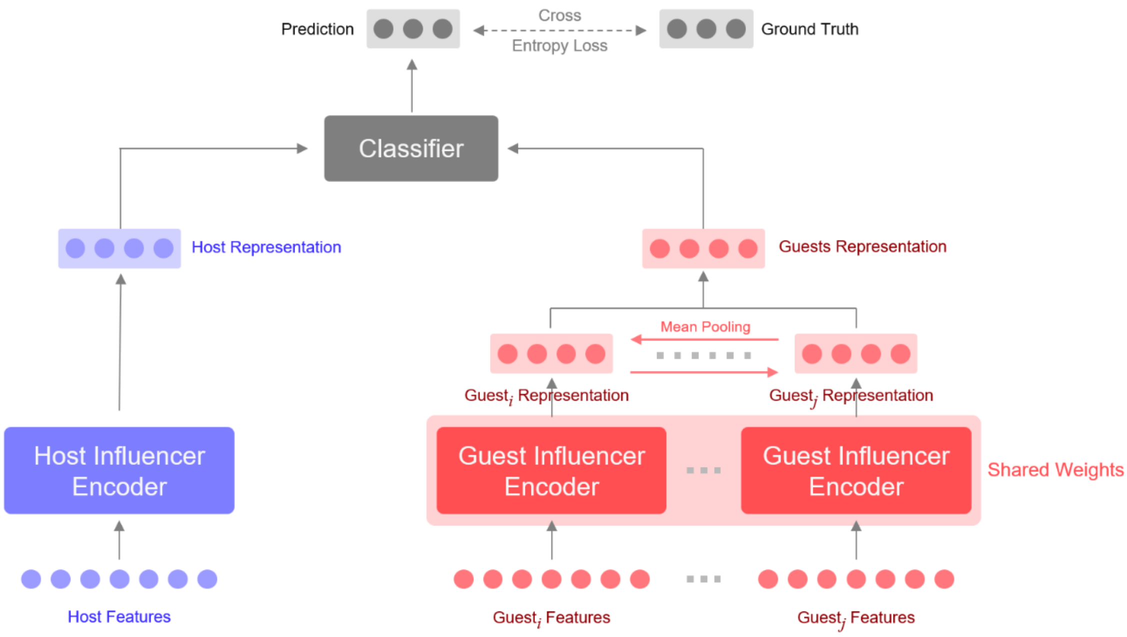

Figure 2: Model Structure of Our Proposed CIIE Method

3.2.2 Influencer Encoder

Now, let us shift our focus to the influencer encoder, specifically the way it handles and preprocesses the raw inputs, as illustrated in the bottom half of Figure 3 , which shows the data flow and transformations applied to obtain the representations for each influencer's channel and their past videos. As mentioned, the influencer encoder plays a crucial role in summarizing all relevant information about the influencer. It shares an identical structure for both hosts and guests, handling similar types of inputs in a consistent manner on both the channel and the video level. Textual features are encoded using Chinese RoBERTa (Cui et al., 2021; Liu et al., 2019), a powerful language

doi:10.6342/NTU202304164

pretrained model that effectively captures the nuances in the textual data. Categorical features are encoded through a learnable embedding layer, allowing the model to learn meaningful representations for each category. To further enhance the model's performance, statistical features are logarithmized, facilitating better feature scaling and handling of large numerical ranges.

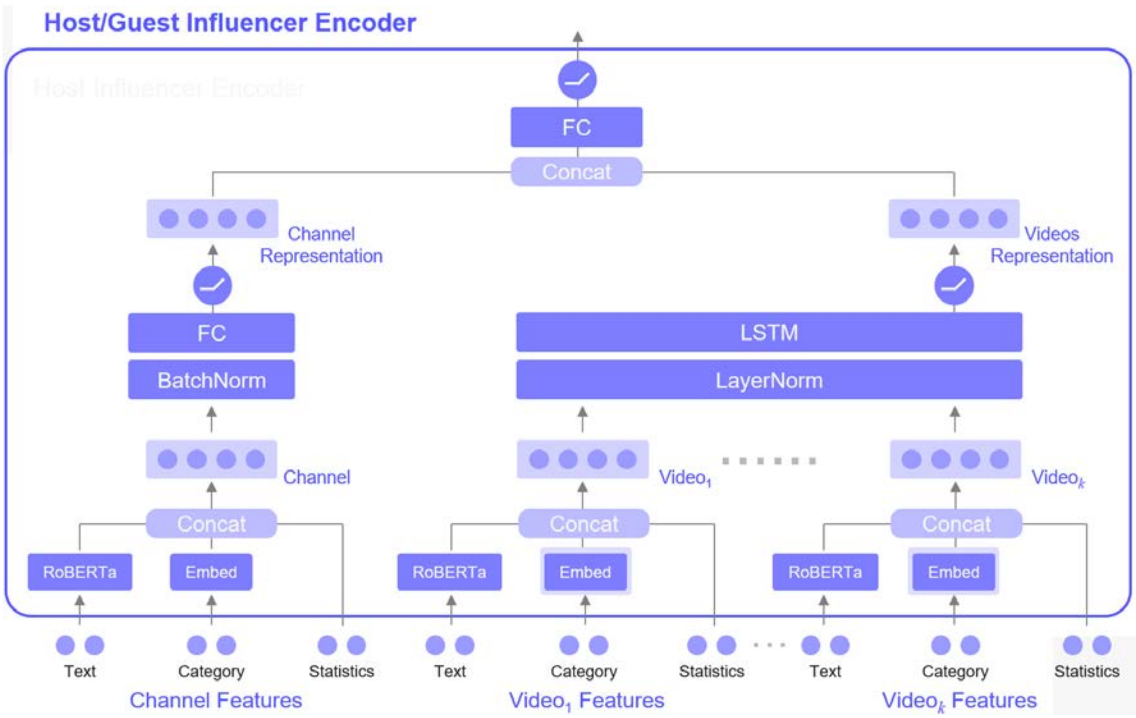

Figure 3: Structure of the Influencer Encoder

It warrants attention that, in our data, we encounter two inherently distinct types of category features: one for channels and the other for videos, each serving a different purpose. The channel's category feature is multi-valued, meaning that a channel can belong to multiple topic categories. Hence, to obtain a fixed-sized embedding for each channel, we take the mean of the embeddings of all the topic categories to which the channel belongs. This approach ensures consistency in the representation of channels, regardless of the number of topic categories they are associated with. On the other hand,

doi:10.6342/NTU202304164

the video's category feature is single-valued, as each video belongs to only one category. For video category embeddings, we directly encode each video's category into a vector without the need for averaging. By handling channel and video category features differently, we can effectively capture the distinct nature of these features and create consistent and meaningful embeddings for both channels and videos in our model.

It is also noteworthy that image features are not considered in our proposed influencer encoder. This is because images are generally encoded into a high-dimensional space, which can make the overall vector space too sparse. For example, the ResNet model encodes each picture into a vector of 1000 dimensions (He et al., 2016). This could potentially affect the model's ability to strike a balance between different types of features, as image features may dominate the data representation. Moreover, incorporating high-dimensional image vectors would introduce a significant number of parameters, leading to a more complex model that typically requires a much larger training dataset and increases the risk of overfitting when the size of training dataset is not large enough. Hence, for this research, the decision is made to omit image features to manage the model complexity and maintain the generalization capability of our proposed model.

After the preprocessing of raw inputs from different channels and videos, each of them has its own representation. Therefore, for each influencer, we obtain a channel representation derived from its channel features. Additionally, we obtain k video representations that respectively stands for the past k videos prior to a given featured video.

Now, let us direct our attention to the middle section of Figure 3. For channel features, they undergo the following layers of processing: (1) a BatchNorm (batch normalization) layer, which normalizes the data to prevent gradient explosion (Ioffe &

doi:10.6342/NTU202304164

Szegedy, 2015; Santurkar et al., 2018), (2) FC (fully-connected) layer, which learns the interactions among the multi-modal inputs and generates a channel representation, and (3) ReLU, an activation function adding nonlinearity to the model (Agarap, 2018). For video features, treated as a sequence, they go through the following layers: (1) LayerNorm (layer normalization), a normalization layer stabilizing sequential data and preventing gradient explosion (Ba, Kiros, & Hinton, 2016), (2) LSTM, a layer that learns sequential information along the inputs and produces a representation for the input video sequence (last output hidden state of the layer) (Hochreiter & Schmidhuber, 1997); and (3) ReLU, an activation function for nonlinear transformation.

The outcome of the above layers is two vectors. The first is the channel representation, which remains consistent across all featured videos and serves as a condensed overview of the influencer's general profile. The second vector is the video representation, which varies for each specific featured video and provides a snapshot of the influencer's current state before entering into a collaboration with others. Together, these two vectors play a crucial role in capturing and summarizing all relevant information about the influencer, enabling the model to make informed predictions of the given featured video. Finally, these two vectors are passed through a fully-collected layer and a ReLU function to summarize all of the information about a given influencer, as shown in the top section of Figure 3.

3.2.3 Mean-pooling Layer

By now, we obtain a host representation and potentially multiple guest representations. Then, before feeding all these representations into the final classifier, a

doi:10.6342/NTU202304164

mean pooling layer is utilized to generate an overall representation for multiple guests. Recall that the shared guest influencer encoder produces a variable number of guest representations. These are passed through the mean pooling layer, which effectively combines and consolidates the information from all guests to arrive at an average representation for all guests. This approach ensures that the model accurately captures the collective influence and impact of all guest influencers on the featured video's performance.

3.2.4 Classifier

To generate the probability distribution for each class, a straightforward classifier is employed. The host representation and the overall guest representation are first concatenated, which is then passed through two fully-connected layers, with a ReLU function in between to introduce nonlinearity, allowing the model to learn and capture the intricate relationships between the host and the guests. Finally, a softmax function is applied to the output, which yields the probabilities for each class based on the given inputs. The structure of the final classifier is illustrated in Figure 4.

doi:10.6342/NTU202304164

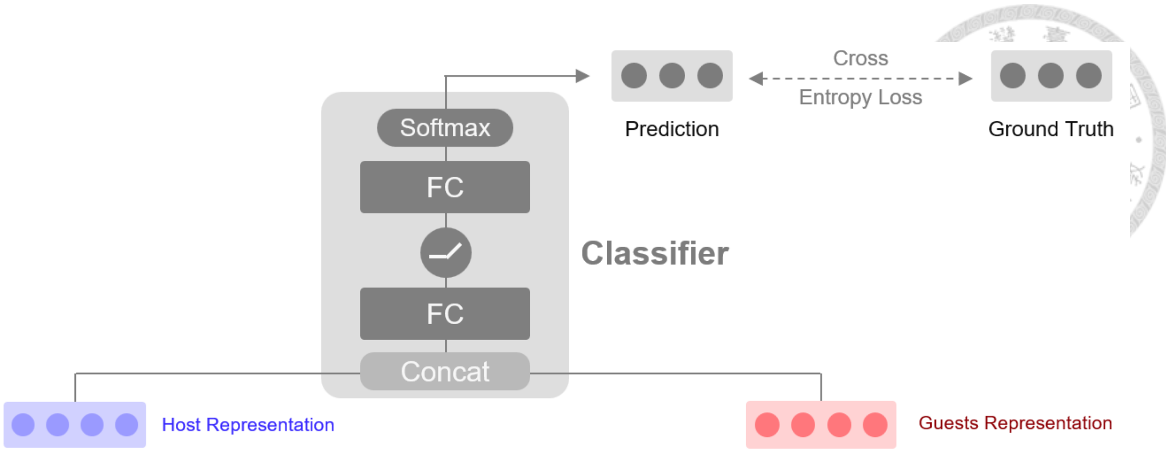

Figure 4: Structure of the Final Classifier

3.2.5 Class Imbalance Method

To address the severe class imbalance problem, we employ class weights as a strategy. The weight w_i assigned to each class i is determined by the reciprocal of its proportion in the training data. The formula is shown below in (3.1), where N_j denotes the number of training data for the j-th class.

$$w_i = \sum_j N_j / N_i \tag{3.1}$$

As a result, higher class weights are assigned to the minority classes, namely, the cannibalized and boosted classes in our research, while a lower class weight is assigned to the majority class, namely, unaffected. This approach helps the model focus more on learning from the imbalanced classes, which are critical for the research focus. For the training process, we utilize a cross-entropy loss function, which is commonly used for multi-class classification tasks, enabling the model to optimize and make accurate predictions for the different classes in our cannibalization prediction problem.

doi:10.6342/NTU202304164

Chapter 4 Empirical Evaluation

4.1 Dataset

4.1.1 Data Collection

To investigate the cannibalization prediction problem, we gathered channel and video data from 166 influencers in Taiwan using the API of YouTube, the largest video sharing site in the world (Zhou et al., 2016). These influencers are called YouTubers. It is important to note that we, by definition, excluded non-influencer channels, such as singers, musicians, brands, movies, cartoons, and media (news, TV, and radio stations), to align with our research objectives.

The data collection process was conducted in three phases. Initially, we gathered data from the top 50 Taiwan YouTubers, as meticulously listed in Wikipedia in the year 2022. This initial set of influencers served as a representative sample of prominent content creators in Taiwan. Moving forward, we expanded the dataset by curating a subset of the most frequently featured guests in the videos of these top YouTubers, as discerned through their video titles. Finally, we repeated the same process to gather a third batch of influencers to expand the dataset further.

As with any endeavor involving data collection from dynamic online platforms, it is essential to acknowledge that the data collected may not perfectly reflect the current scenario, as metrics like video views and subscriber counts can change over time due to the dynamic nature of influencer channels and audience engagement patterns. Despite this, our dataset provides valuable insights into the cannibalization phenomenon and its implications within the context of featured videos by social media influencers in Taiwan.

doi:10.6342/NTU202304164

4.1.2 Overview of Our Dataset

Our dataset is structured at the video level and comprises two main components: channel features and video features. Both of these components are further divided into four types of information, namely, text, category, statistics, and images. The channel features consist of textual, categorical, statistical, and image-related information. The text features include the channel title, channel description, and channel keywords. The category features capture the topic categories associated with the channel. The statistical features encompass metrics such as the view count, subscriber count, video count, and published date, offering quantitative data on the channel's popularity and activity. Lastly, the image features include the channel thumbnail and channel banner, visually representing the influencer's branding and style.

On the other hand, the video features also encompass text, category, statistical, and image-related information. The text features of the video comprise the video title, video description, and video tags, which provide textual context and information about the content of the specific video. The category feature is the video category, which can be configured by influencers and typically includes labels such as "Gaming", "Entertainment", "Education", and many others. The statistical features of the video include its duration, like count, view count, and comment count. Lastly, the image features of the video consist of the video thumbnail.

Table 2 presents the summary statistics of our dataset, focusing on the channel level. Note that the process of identification of featured videos to generate statistics related to the average number of featured videos will be detailed later in Section 4.1.4. Further, it

doi:10.6342/NTU202304164

is important to note that during the data collection process, we might gather data from some micro-influencers with smaller size of followers. While the majority of the influencers are significant players in the industry with substantial subscriber counts and video views, a few micro-influencers were included due to their past guest appearances on larger channels. For example, the smallest one has only 11,300 subscribers and 7 uploaded videos, none of which is a featured video. Despite this inclusion, we ensured rigorous data collection process to maintain the dataset's quality and relevance to our research objectives.

Table 2: Summary Statistics of Our Dataset

	Min	Max	Avg	Std
Number of subscribers	11,300	5,710,000	940,570	899,158
Total video views	109,271	2,328,844,865	268,593,094	372,036,752
Average views per video	0.83	17,766.59	2,049.08	2,838.24
Average featured videos	0	874	90.88	126.38

4.1.3 Featured Video Identification

To identify a featured video and its guests, we employ specific criteria to ensure accurate detection of collaborations between YouTubers. First, we look for keywords such as "feat", "ft", or "@," present in the video title, which often indicate a collaborative effort between the host and guest YouTubers. This helps us identify potential featured videos that involve influencer collaborations.

Second, we analyze the video title to check if it includes the name of at least one YouTuber other than the host. The presence of another YouTuber's name suggests the

doi:10.6342/NTU202304164

involvement of a guest channel in the featured video. This step allows us to distinguish between regular videos and those that feature collaborations. Finally, the guest channel(s) mentioned in the video title must be one of the 166 YouTubers we collected during our data collection process. This is because our dataset is limited to the influencers we have identified and gathered information for. Since we cannot identify a YouTuber who is not part of our dataset, ensuring that the guest channels are among the 166 YouTubers helps maintain the integrity and scope of our research.

It is important to mention that this process involves manual research and compilation of each YouTuber's different names and nicknames. Moreover, each YouTuber is associated with a unique and short channel string identifier, known as the username handle, provided by YouTube. This information assists in accurately identifying and verifying the presence of guest YouTubers in featured videos, enabling us to analyze the impact of collaborations in our research work effectively.

4.1.4 Featured Video Identification

To label each featured video as cannibalized, boosted, or unaffected, we employed a sliding window approach with a size of (k + 1 + k) to analyze the videos of each YouTuber. Regarding the size of the sliding window, the first k represents the k videos prior to the featured video, followed by the number 1 denoting the featured video itself, and finally, the second k representing the k videos after the featured video. The sliding window scans through all videos in chronological order for each YouTuber. When the sliding window positions a featured video at its center, we want to compare the average views of its previous k videos and its next k videos. However, before calculating these

doi:10.6342/NTU202304164

averages, we conduct a simple extreme removal process. We remove the videos with the highest and lowest views from both the previous *k* videos and the next *k* videos within the same window. This ensures that the labeling of any featured video is not affected by extreme values.

Finally, if the average view of the featured video's previous *k* videos drops or rises by θ% after the featured video, we label this featured video as "cannibalized" or "boosted," respectively. Otherwise, the video is labeled as "unaffected." This systematic approach helps us identify and classify the impact of each featured video accurately, taking into account the surrounding video views and avoiding bias from extreme values.

To illustrate, let us take a look at Figure 5 and Figure 6, which represent featured videos labeled "cannibalized" and "boosted", respectively. Similar to Figure 1, each vertical bar represents a video and the height of the bars represents their respective video views. The red bar, located at the central position, corresponds to a featured video. During the calculation of average views, the grey bars, representing videos with minimum or maximum views, are removed to ensure more accurate averages. The blue bars indicate the videos used to calculate the averages. The averages are represented in horizontal red dashed lines, with their exact numbers annotated slightly above. Finally, the percentage change of the two averages is denoted around the middle of the figures, with positive numbers in green and negative numbers in red. This visualization allows us to observe the patterns in views and identify potential cannibalization or boosting effects around the featured video.

doi:10.6342/NTU202304164

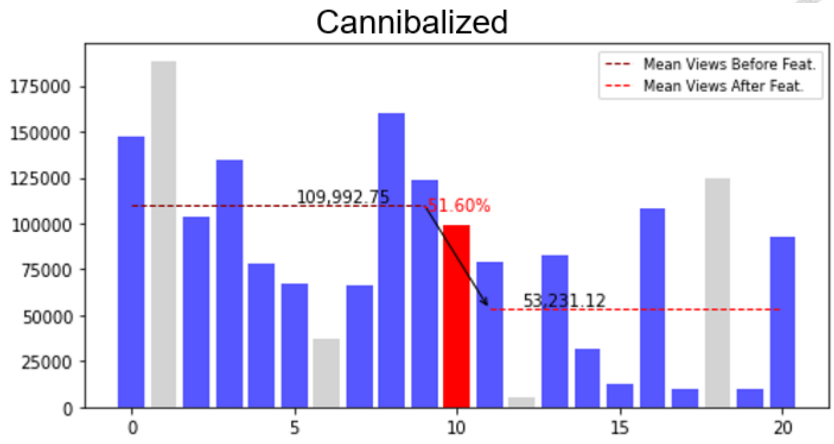

Figure 5: An Example of A Featured Video Labeled as “Cannibalized”

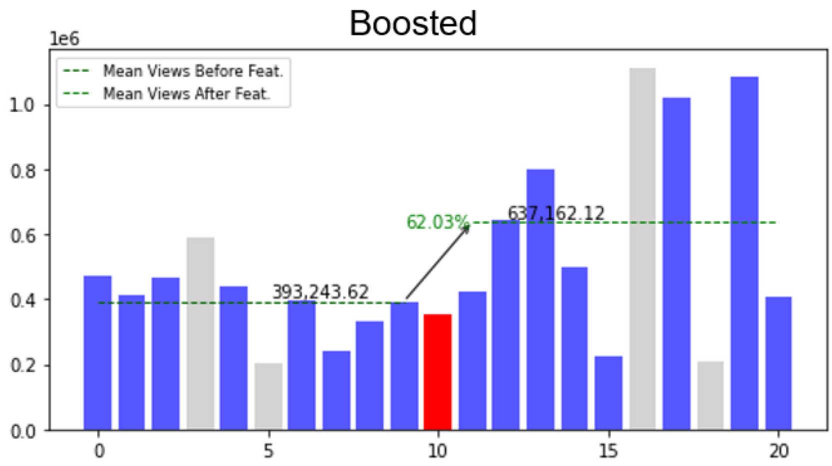

Figure 6: An Example of a Featured Video Labeled as “Boosted”

For the featured video in Figure 5, the average views following the featured video dropped a significant 51.6% as compared to the average views before. Due to this substantial decline, the video is labeled as “cannibalized.” This label indicates that the featured video impacts the views of subsequent videos, suggesting the featured video creates cannibalization effect on the host influencer. For the featured video in Figure 6, the average views following the featured video experienced a significant increase of 62.03% as compared to the average views before. Due to this remarkable increase, the

doi:10.6342/NTU202304164

video is labeled as “boosted.” This label indicates that the video’s performance positively impacts the views of subsequent videos, showing a boosting effect.

4.1.5 Featured Video Filtering Process

First, we perform a filtering step based on proximity. Specifically, we remove featured videos that are too close to each other within the same window, unless all of them are labeled as “unaffected” videos. This decision is made to avoid potential ambiguity in determining which featured videos are responsible for significant rises or drops in average views when they are clustered together.

When multiple featured videos are within the same window, it becomes challenging to attribute the observed changes in average views to a specific video. For example, in Figure 7, it is hard to determine whether the general decline of views is caused by which featured videos (colored in red). Along similar lines, in Figure 8, the influence of the two boosting featured videos (colored in green) on the upward trend in video views cannot be accurately discerned because of their proximity to each other within the window. To address this, we opt to remove all the featured videos in that same window, including the central one and its neighboring featured videos. By doing so, we can maintain the clarity and reliability of our predictions, ensuring that the impact of each featured video on subsequent views is accurately captured and attributed.

doi:10.6342/NTU202304164

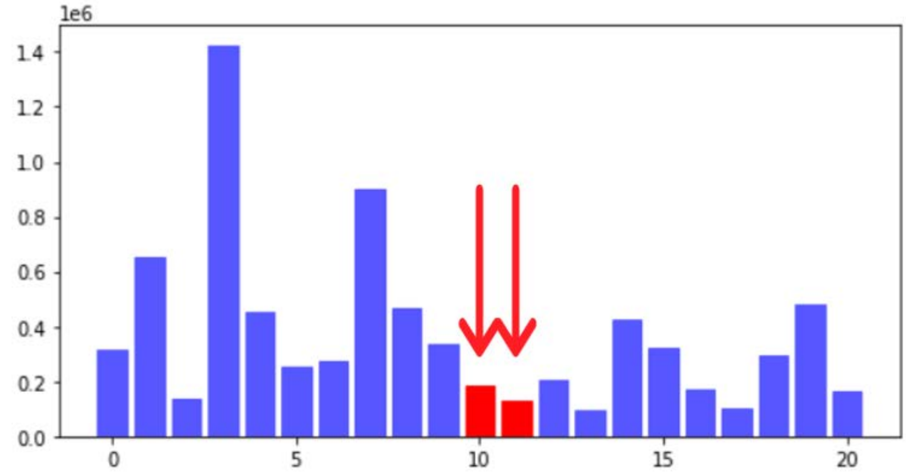

Figure 7: An Example of Two "Cannibalized" Featured Video Too Close Together

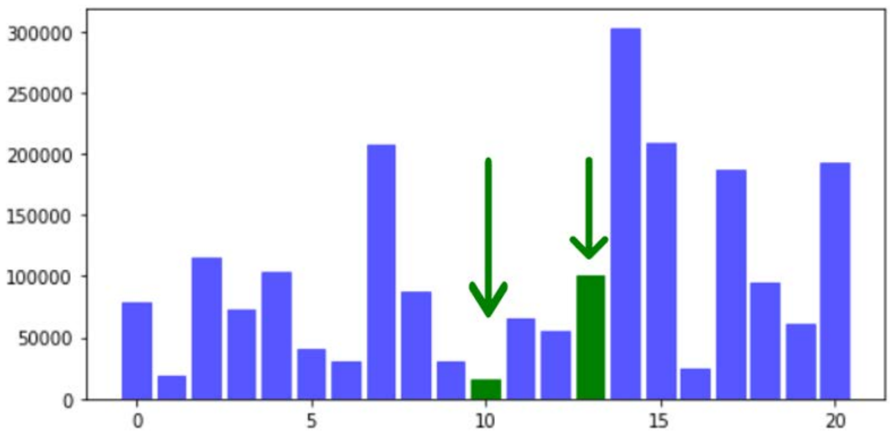

Figure 8: An Example of Two "Boosting" Featured Video Too Close Together

To clarify further, featured videos within the same window are retained only when all of them are labeled as "unaffected" because they do not exhibit significant changes in views. In this case, there is no need for removal as these videos do not contribute to any notable rise or drop in average views, and they are not likely to confound the analysis.

Notably, when filtering out featured videos in close proximity, we opted to disregard the potential effect of featured videos on the views of the subsequent *k* videos of the guests. This decision was made due to the fact that the featured video is always uploaded

doi:10.6342/NTU202304164

on the host's channel. As a result, the viewers of the guest influencers cannot see or may not have direct access to the host's videos, leading to a reduced likelihood of the featured video significantly affecting the viewership of the guests' subsequent content. This consideration guided the decision to not factor in the influence of the host's uploaded featured videos on the viewership of the guests' subsequent videos.

As a final step, we filter out featured videos that have more than 5 guests. This decision is based on the observation that videos with a large number of guests are relatively uncommon and are predominantly categorized as "unaffected." As our research primarily focuses on identifying and understanding the cannibalization and boosting effects, which are often associated with specific guest influencers, these videos with multiple guests hold little relevance to our study.

By excluding these videos, we remove a total of 180 entries, approximately 3.2% of the total featured videos in the dataset. This filtering process ensures that our model can focus on the more relevant videos and reduce unnecessary complexity while maintaining the accuracy and effectiveness of our predictions. Table 3 presents the distribution of labels for each class, grouped by the number of guests in the featured videos, before the removal process.

doi:10.6342/NTU202304164

Table 3: Featured Video Label Distribution Grouped by Number of Guests

	Number of labels		
#guests	boosted	cannibalized	unaffected
1	210	111	3,131
2	65	40	941
3	23	18	494
4	21	7	261
5	11	5	109
6	4	1	78
7	1	2	43
8	5	0	31
9	1	0	3
10	1	0	4
11	0	0	1
12	0	1	2
14	1	0	1

After applying the filtering and preprocessing steps, the final dataset consists of 5,447 featured videos. For the cannibalization prediction task, we set the parameters as k = 10 (representing the number of videos before and after the featured video for average view comparison) and θ % = 40% (the threshold for identifying cannibalization or boosting effects). Note that the average percentage change in views before and after these 5,447 featured videos is 18.5%, meaning that featuring other influencers in a video is indeed generally a good strategy and that dropping 40% in views, defined as cannibalization in this work, is certainly a critical problem. With these settings, we are left with 146 YouTubers who have at least one featured video that meets the criteria and is not excluded. The summary statistics of these channels are presented in Table 4.

doi:10.6342/NTU202304164

Table 4: Summary Statistics of Channels with Featured Videos (after Filtering)

Channel Count	146
Avg #Subscribers	851,357.53
Avg Total Views	245,371,565.84
Avg Views per Video	298,434.00
Avg #FeaturedVideos	38.22

On the other hand, as shown in Table 5, the dataset undoubtedly suffers from class imbalance. Particularly, the "cannibalized" class, which is the main target of our study, accounts for only 3.32% of the total featured videos. The skewed distribution of the three classes highlights the challenge of dealing with class imbalance and underscores the need for careful handling and evaluation of the predictive model.

Table 5: Class Distribution in Featured Videos (After Filtering)

	Count	**Percentage**
Unaffected	4,936	90.62%
Boosted	330	6.06%
Cannibalized	181	3.32%

For each video in the obtained set of featured videos, we extract the following predictors from our original dataset: (1) channel features of the host channel, (2) channel features of all guest channels, (3) video features of host's past *k* videos prior to the given featured video, and (4) video features of guests' past *k* videos prior to the given featured video. With this rich set of predictor features, our objective is to classify each featured video into one of three classes: cannibalized, boosted, or unaffected.

doi:10.6342/NTU202304164

4.2 Evaluation Procedure

4.2.1 Experimental Setup

For our experimental setup, we employ stratified 5-fold cross-validation to evaluate the performance of our model. Before this, we separate a test set consisting of 897 data points for final evaluation. The remaining 4,550 data points are then divided into 5 folds while ensuring that the proportions of each class remain consistent across the training, validation, and testing sets. In each fold, the training set consists of 3,640 data points, and the validation set has 910 data points.

To assess the performance of our model, we rely on precision, recall, and F1 scores as our key evaluation metrics. Our primary focus is on maximizing the precision and recall scores for the minority classes. In particular, we consider the performance on the "cannibalized" class to be of utmost importance, even more than the "boosted" class.

4.2.2 Model Configuration

For our experiments, we have set the batch size to 512, which determines the number of samples processed in each training iteration. A learning rate of 1e-5 has been chosen to control the step size in updating the model weights during training. The number of epochs is set to 2000, indicating the number of times the entire dataset is passed through the model during training. To prevent overfitting and improve efficiency, we have implemented early stopping with a patience of 200 epochs. Early stopping stops the training process when the model's performance on the validation set does not improve for a specified number of epochs, preventing unnecessary computation. The optimizer

doi:10.6342/NTU202304164

used is Adam, a popular algorithm that adapts the learning rate during training to achieve faster convergence. Lastly, we set the random seed to 42 to ensure reproducibility of our results. The hyperparameters of our model are presented in Table 6.

Table 6: Hyperparameters of Our Proposed CIIE Model

Hyperparameter	**Value**
Batch size	512
Learning rate	1e-5
Number of epochs	2000
Early stopping	200
Optimizer	Adam
Random seed	42

The model's hidden dimensions are configured as follows: the embedding size for categorical features is set to 10, providing a compact representation for category-related information. The hidden dimension for channel representation is set to 256, capturing essential characteristics about the host influencer's channel. Similarly, the hidden dimension for video representation is set to 256, encoding relevant information from the videos. For the influencer representation, a hidden dimension of 128 is utilized, summarizing the overall profile of the influencer based on their channel and video features. The hidden dimension before the output layer is set to 64, acting as a bottleneck layer to further distill the learned information. Finally, the output dimension is set to 3, representing the three classes for classification: cannibalized, boosted, and unaffected. These carefully selected hidden dimensions are crucial for the model's performance and ability to capture meaningful patterns in the data.

doi:10.6342/NTU202304164

Table 7: Hidden Dimensions of Our Proposed Model

Hidden Dimension	**Value**
Embedding size (of categorical features)	10
Hidden dimension for channel representation	256
Hidden dimension for video representation	256
Hidden dimension for influencer representation	128
Hidden dimension before output layer	64
Output dimension	3

4.2.2 Performance Benchmarks

We employ the following benchmarks for evaluating our model's performance. The first benchmark is the Prior Probabilistic Model (PPM), which takes into account the probability of each class when a guest influencer collaborates with the host. By calculating the prior probability for each class, the model predicts the class with the highest probability as the label for the featured video when a single-guest case is considered. This approach assumes that certain guests might have a higher chance of causing cannibalization or boosting effects.

We also evaluated our model against the Constrained Prior Probabilistic Model (CPPM). This benchmark restricts the proportion of the minority classes, namely cannibalized and boosted, to a predetermined level α. By doing so, the CPPM controls the distribution of these classes and aims to improve the prediction accuracy for the minority classes.

The final benchmark is the Random Guess Model (RGM), which randomly assigns labels to the featured videos based on a pre-determined class distribution. This benchmark

doi:10.6342/NTU202304164

serves as a baseline comparison to assess the predictive power of our model compared to random guessing.

4.2.3 Benchmark 1: Prior Probabilistic Model (PPM)

The training process for the Prior Probabilistic Model (PPM) involves the following steps:

1. Calculate Probabilities: For each channel in the training data, the model calculates the probability of each class (cannibalized, boosted, unaffected) when the channel is a guest to another channel. This step involves analyzing historical data and guest appearances to compute the likelihood of each class for a specific channel.

2. Normalize Probabilities: After obtaining the probabilities, they are normalized to Z-scores with respect to each class across all channels. Normalization helps to standardize the probabilities and bring them to a common scale, making comparisons between different channels more meaningful.

3. Assign Classes: Based on the Z-scores, the PPM assigns the class with the highest Z-score to the channel. The class with the highest likelihood becomes the predicted label for the featured video when the channel acts as a guest.

4. Handle Ties: In cases where there is a tie in the Z-scores, indicating that multiple classes have similar probabilities, the PPM randomly selects one of the tied classes to be the final assignment. This ensures a fair and unbiased decision when there are equal probabilities for different classes.

An illustration of the training process is provided in Table 8. In the table, each row represents a guest channel, and the columns indicate the Z-scores for each class. Z-scores

doi:10.6342/NTU202304164

standardize the probabilities of different classes, enabling fair comparisons. The prediction column indicates the final class assigned to each guest channel based on the highest Z-score. For example, “Guest 1” has the highest Z-score for the “cannibalized” class, making it the predicted class for that guest channel.

The PPM applies this process to all guest channels in the training data to determine the most likely class labels for featured videos when specific channels act as guests. The model’s ability to calculate Z-scores aids in making informed predictions about potential cannibalization or boosting effects of featured videos based on historical guest appearances.

Table 8: An Illustration of PPM’s Training Process

	Z-scores			
	boosted	cannibalized	unaffected	**prediction**
Guest 1	-0.568167	-0.396693	-0.628325	cannibalized
Guest 2	-0.568167	-0.626123	-0.628325	boosted
Guest 3	-0.363838	-0.626123	-0.671523	boosted
Guest 4	-0.772496	-0.855554	-0.73632	unaffected
Guest 5	-0.159509	-0.396693	-0.282739	boosted
...	...	...	...	...
Guest 142	-0.363838	-0.855554	-0.585126	boosted
Guest 143	-0.363838	-0.626123	-0.239541	unaffected
Guest 144	NaN	NaN	NaN	unaffected
Guest 145	NaN	NaN	NaN	unaffected
Guest 146	NaN	NaN	NaN	unaffected

The inference process for the Prior Probabilistic Model (PPM) involves determining the class label for a given featured video by considering the class assignments of its guest channels. The steps are as follows: (1) Given a featured video, the model looks up the corresponding class assignments of its guest channels in the training data. These class

doi:10.6342/NTU202304164

assignments were determined during the training phase based on the highest Z-scores for each guest channel. (2) If a guest channel does not exist or did not appear as a guest in the training data, the model assigns it the majority "unaffected" label. (3) The final class assignment for the featured video is determined by the majority vote, which is the class with the most occurrences among the guests. In case of a tie, where two or more classes have the same number of occurrences, the model randomly selects one of the majority classes to be the final prediction.

4.2.4 Benchmark 2: Constrained Prior Probabilistic Model (CPPM)

The training process for the Constrained Prior Probabilistic Model (CPPM) adds additional steps to that of PPM, which involves the following: (1) Define a minority percentage α, which represents the maximum allowable proportion of channels assigned to the "cannibalized" or "boosted" class. For instance, if α is set to 10%, it means that only the top 10% of channels with the highest Z-scores in the "cannibalized" or "boosted" label will be allowed to be assigned to these classes. (2) Run the training process of the Prior Probabilistic Model (PPM), as described earlier, to obtain initial class assignments for all channels. (3) Impose the constraint on the percentage of "cannibalizing" and "boosting" channels. For channels assigned to the "cannibalized" class, the model identifies those whose Z-scores in the "cannibalized" label are not in the top α percent and reassigns them to the "unaffected" class. The same process applies to channels assigned to the "boosted" class.

By reassigning some channels from "cannibalized" or "boosted" to "unaffected," the CPPM ensures that the percentage of channels in the "cannibalized" and "boosted"

doi:10.6342/NTU202304164

classes does not exceed the pre-determined proportion α. This constraint maintains a lower percentage of the minority classes and prevents them from dominating the model's predictions. During inference, CPPM uses the same mechanism as PPM to make predictions based on the class assignments of guest channels.

4.2.5 Benchmark 3: Random Guess Model (RGM)

The Random Guess Model (RGM) is trained by randomly assigning labels to each channel in the training data according to a pre-determined probability distribution of the three classes (cannibalized, boosted, and unaffected). For instance, the model might be set with the probability distribution (10%, 10%, 80%) for the three classes, which means there is a 10% chance of assigning the "cannibalized" label, a 10% chance of assigning the "boosted" label, and an 80% chance of assigning the "unaffected" label to a channel during training.

During inference, the RGM also uses the same process as PPM to make predictions for each featured video. It looks up the corresponding class assignments to its guest channels, and if a guest channel does not exist or appear in the training data, it is assigned to the "unaffected" label. The model then takes a majority vote on these assignments and uses the final result as the prediction for the featured video's class.

To calculate the evaluation metrics, the training and inference process is run 1000 times for each fold, and the averages of these results are taken. This helps to account for the randomness introduced during the training of the RGM and provides a more stable estimate of its performance.

doi:10.6342/NTU202304164

4.3 Evaluation Results

4.3.1 Experiment Results

The comprehensive comparison of the performance of the benchmarks and our model is presented in Table 9. In a general view, our model demonstrates superior overall performance in terms of precision, recall, and F1-score compared to other benchmarks. Particularly, our model outperforms other methods on the minority classes, namely “cannibalized” and “boosted,” except for the recall on the “boosted” class. However, our model’s recall on the “boosted” class is still comparable to the best one achieved by the Prior Probabilistic Model (PPM), which tends to favor predicting minority classes, as we will discuss later.

Moreover, our model’s results on the macro average of the metrics are also better than those of other benchmarks. However, since the macro average considers the outcome of the unaffected class as well, it is not our primary focus of interest in this study. Instead, we place more emphasis on the performance of our model on the minority classes, which are of greater significance in the context of cannibalization prediction.

Now, let’s dive deeper into the results for further findings. In terms of minority classes, our model exhibits superior performance in the “cannibalized” and “boosted” classes concerning precision, recall, and F1-score compared to other benchmarks. The precision values for the minority classes are notably higher than those achieved by other methods. This indicates that our model effectively captures critical information that aids in detecting both cannibalization and boosting effects in potential featured videos.

Furthermore, the higher recall in these classes suggests that our model successfully identifies more true positive instances, thereby demonstrating its capability to capture

doi:10.6342/NTU202304164

relevant patterns related to cannibalization and boosting. These results validate the effectiveness of our efforts to address the class imbalance problem, as our model delivers promising outcomes for the minority classes, which are the primary focus of our research.

On the other hand, PPM ranks as the second-best model in terms of overall metrics across all classes, indicating that the Z-scores of each class for each channel provide valuable information for predicting cannibalization effects. It outperforms the RGM baseline in nearly all metrics, making it a comparable benchmark for evaluation. Notably, PPM exhibits significantly higher recall for minority classes and lower recall for the majority class, suggesting its preference for predicting minority classes.

RGM showcases the highest recall and F1-scores in the majority class, which can be attributed to the pre-determined probability of sampling the “unaffected” label at 90%. Consequently, its recall also hovers around 90%, specifically 91.39%. The exceptionally high recall score has a consequential effect of boosting the F1-score. However, it is important to recognize that our research primarily centers around the minority classes, and thus, we can safely disregard concerns pertaining to the outstanding recall and F1-score of the majority class, as it does not align with our primary objectives.

CPPM falls short compared to PPM in every metric. Its lower recall scores make sense because as we set a lower value for α, the recall of the minority classes naturally drops, while that of the majority class increases, as we explicitly set out to predict fewer minority labels and more majority labels. However, the drop in precision for all classes in CPPM is not expected. This might be attributed to the fact that the Z-scores of the minority classes do not account for their frequency. For example, considering a scenario where a channel is a guest only once, the corresponding Z-score of the label “cannibalized”

doi:10.6342/NTU202304164

is bound to be high. However, this isolated occurrence may not be predictive of future collaborations, leading to overall lower precision.

CPPM, as comparing to RGM, achieves higher recall, specifically when considering only the minority classes. This indicates that CPPM has a better ability to capture more true “cannibalized” and “boosted” labels than RGM. Meanwhile, there is no notable difference in precision between the two models. This implies that both CPPM and RGM are equally skilled at avoiding false positive predictions. Overall, while CPPM falls short compared to PPM, it still demonstrates an advantage over RGM in terms of recall.

Table 9: Evaluation Results of Benchmarks and Our Proposed CIIE Model

	Cannibalized			**Boosted**			**Unaffected**			**Macro Avg**		
	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score
RGM (10%-10%-80%)	0.0862	0.0871	0.0830	0.0902	0.0860	0.0841	0.8232	0.8219	0.8217	0.3332	0.3317	0.3296
RGM (5%-5%-90%)	0.0879	0.0427	0.0536	0.0903	0.0413	0.0528	0.8237	**0.9139**	**0.8661**	0.3340	0.3326	0.3242
PPM	0.0957	0.3929	0.1538	0.1238	**0.4533**	0.1942	**0.8720**	0.3348	0.4828	0.3638	0.3937	0.2769
CPPM (α = 10%)	0.0890	0.1905	0.1209	0.1183	0.3000	0.1695	0.8389	0.5963	0.6968	0.3487	0.3623	0.3291
CPPM (α = 5%)	0.0777	0.1071	0.0893	0.0813	0.1222	0.0974	0.8222	0.7444	0.7811	0.3271	0.3246	0.3226
Our Model	**0.1471**	**0.4833**	**0.2246**	**0.1353**	0.4111	**0.2014**	0.8707	0.4647	0.6000	**0.3844**	**0.4531**	**0.3420**

4.3.2 Ablation Test 1: Effects of Different Inputs on Model Performance

Removing any type of inputs from the model results in a decline in the overall performance, as shown in Table 10. For the “cannibalized” class, we observe that statistics features have the most influence on both precision and recall scores, while text features show the least impact. On the other hand, for the “boosted” class, category features play a significant role in both precision and recall scores, whereas statistics features have the least impact.

doi:10.6342/NTU202304164

The metrics for the majority class are not the best, but they show little difference compared to other scenarios. However, since our research primarily focuses on the minority classes, this should not be a major concern.

Table 10: Ablation Results on the Effects of Different Model Inputs

	Cannibalized			Boosted			Unaffected			Macro Avg		
	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score
Our Model w/o Text	0.1351	0.4786	0.2107	0.1338	0.4089	0.2013	0.8638	0.4412	0.5837	0.3776	0.4429	0.3319
Our Model w/o Cats	0.1313	0.4286	0.1981	0.1304	0.3778	0.1915	**0.8714**	**0.4760**	**0.6055**	0.3777	0.4274	0.3317
Our Model w/o Stats	0.1117	0.4524	0.1786	**0.1354**	0.4000	**0.2014**	0.8813	0.4137	0.5608	0.3761	0.4220	0.3136
Our Model	**0.1471**	**0.4833**	**0.2246**	0.1353	**0.4111**	**0.2014**	0.8707	0.4647	0.6000	**0.3844**	**0.4531**	**0.3420**

4.3.3 Ablation Test 2: Effects of Image Inputs on Model Performance

In this ablation study, we encode the images using the ResNet model into 1000-dimensional vectors. These are concatenated with other encoded features, such as textual representations encoded by Chinese RoBERTa, learnable category embeddings, and statistics features, and then processed in the same way as others. The results are shown in Table 11. The inclusion of image inputs, such as thumbnails and banners, negatively impacts the overall performance of the model. Specifically, there is a notable decline in all metrics related to the minority classes. As discussed earlier, we believe that the addition of high-dimensional image vectors makes the vector space too sparse and model struggle to balance all types of features. Moreover, the increased complexity from adding excessive parameters significantly slows down the training process and could lead to overfitting.

Interestingly, removing image inputs only leads to a minor improvement in the recall and F1-score of the majority class. However, since our research primarily focuses on the

doi:10.6342/NTU202304164

minority classes, this improvement in the majority class is not our main concern. Therefore, the negative impact of image inputs on the overall model performance outweighs any marginal benefits they may bring to the majority class.

Table 11: Ablation Results on the Effect of Image Inputs

	Cannibalized			Boosted			Unaffected			Macro Avg		
	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score
Our Model w/ Image	0.1159	0.4095	0.1767	**0.1186**	0.2956	0.1687	0.8611	**0.4875**	**0.6154**	0.3652	0.3975	0.3203
Our Model	**0.1471**	**0.4833**	**0.2246**	0.1353	**0.4111**	**0.2014**	**0.8707**	0.4647	0.6000	**0.3844**	**0.4531**	**0.3420**

4.3.4 Effect of Class Imbalance Methods on Classification Effectiveness

As demonstrated in Table 12, the use of class weights has proven to be the most effective method in addressing the class imbalance problem, resulting in superior overall model performance. Unlike other class imbalance methods, class weights strike a better balance between precision and recall across all classes.

When no measures are taken to handle class imbalance, the model tends to perform poorly, simply classifying all samples into the majority class. This highlights the necessity of employing appropriate techniques to tackle the class imbalance issue. While methods like ADASYN (He et al., 2008) and SMOTE (Chawla et al., 2002) show slight improvements in precision for the cannibalized and boosted labels, respectively, they come at a significant cost of remarkably low recall in the minority classes. This indicates that these methods may prioritize the positive identification of certain classes, sacrificing the ability to capture important patterns related to cannibalization and boosting.

Moreover, the Borderline SMOTE (Han, Wang, & Mao, 2005) method further worsens the model's performance across all metrics, indicating its limited effectiveness

doi:10.6342/NTU202304164

in this particular context. As a result, class weights remain the most viable and effective option for handling class imbalance and achieving balanced model performance across all classes.

Table 12: Effect of Class Imbalance Methods on Classification Effectiveness

	Cannibalized			**Boosted**			**Unaffected**			**Macro Avg**		
	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score
No adjustment	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.8242	1.0000	**0.9037**	0.2747	0.3333	0.3012
ADASYN	**0.1664**	0.0786	0.1060	0.1439	0.0733	0.0965	0.8312	**0.9172**	0.8720	0.3805	0.3564	0.3582
SMOTE	0.1628	0.0881	0.1101	**0.1537**	0.0756	0.0996	0.8295	0.9130	0.8691	0.3820	0.3589	**0.3596**
BorderlineSMOTE	0.1317	0.0714	0.0912	0.1229	0.0667	0.0859	0.8274	0.9044	0.8642	0.3606	0.3475	0.3471
Class Weights	0.1491	**0.4405**	**0.2215**	0.1457	**0.4733**	**0.2221**	**0.8755**	0.4775	0.6151	**0.3901**	**0.4638**	0.3529

doi:10.6342/NTU202304164

Chapter 5 Conclusions

This research has made significant contributions in several key aspects. First, we defined a novel research task, focusing on predicting cannibalization effects in the context of featured videos by social media influencers. Through data analysis, evidence indicates that cannibalization indeed occurs in approximately 3.3% of the featured videos investigated. This finding shed light on the dynamics of influencer collaborations and their potential impact on video performance.

Second, we propose a deep neural network predictive model CIIE to address the cannibalization prediction problem. The model incorporates influencer encoders, which effectively capture essential information about both the host and guest influencers, as well as their content. By leveraging diverse data sources, the model achieves a comprehensive representation of the influencers and their collaborative efforts.

Furthermore, the proposed Cannibalization Identification with Influencer Encoders (CIIE) model demonstrates superior performance when compared to various benchmark methods, including the Prior Probabilistic Model (PPM), Constrained Prior Probabilistic Model (CPPM), and Random Guess Model (RGM). Our model consistently outperforms these benchmarks in terms of precision, recall, and F1-score, particularly for the minority classes of cannibalized and boosted videos.

Overall, this work presents a comprehensive analysis of cannibalization among social media influencers and provides valuable insights into the development of an effective predictive model. Our work also sheds light on several aspects that can be improved in future research. First, expanding the dataset by collecting a larger and more diverse sample of featured videos would enhance the comprehensiveness of the study.

doi:10.6342/NTU202304164

Currently, the remaining featured videos after filtering stand at 5,447, but obtaining more data could provide richer insights into cannibalization effects.

Second, the scope of the data is restricted to Taiwanese YouTubers, thereby limiting its generalizability to influencers on other platforms or in different countries. Collecting additional datasets from different platforms or incorporating YouTubers from different countries represents an interesting research direction. Third, the operational definition of cannibalization in our context uses only video views. However, incorporating user behaviors such as likes and comments can better gauge viewers' preferences and refine the definition.

Fourth, considering the content of the videos itself as a factor in cannibalization analysis could yield valuable insights. Currently, it is disregarded in this preliminary work because the model complexity would be too high given the limited data at hand. However, addressing these challenges and leveraging the content information could offer a better performance in the prediction of cannibalization effects. Lastly, information about the featured video is not employed in this work due to the assumption that our model's application occurs before any collaboration has taken place. However, in situations where two influencers are exploring a potential partnership, leveraging metadata like video title and category could be deemed reasonable, which opens up a new research direction. For instance, they could brainstorm various video titles or categories and assess them using the model to identify the combination that yields the lowest probability of cannibalization or the highest probability of boosting effects.

doi:10.6342/NTU202304164

References

Agarap, A. F. (2018). Deep learning using rectified linear units (relu). arXiv preprint arXiv:1803.08375.

Ahn, J., Kim, A., & Sung, Y. (2021). The effects of sensory fit on consumer evaluations of co-branding. *Leveraged Marketing Communications* (pp. 42-59).

Alalwan, A. A., Rana, N. P., Algharabat, R., & Tarhini, A. (2016). A systematic review of extant literature in social media in the marketing perspective. *Social Media: The Good, the Bad, and the Ugly* (pp. 79-89). https://doi.org/10.1007/978-3-319-45234-0_8

Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization. arXiv preprint arXiv:1607.06450.

Bouten, L. M., Snelders, D., & Hultink, E. J. (2011). The impact of fit measures on the consumer evaluation of new co-branded products. *Journal of Product Innovation Management*, *28*(4), 455-469.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16(1), 321-357.

Copulsky, W. (1976). Cannibalism in the marketplace. *Journal of Marketing*, *40*(4), 103-105.

Cui, Y., Che, W., Liu, T., Qin, B., & Yang, Z. (2021). Pre-training with whole word masking for Chinese BERT. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 3504-3514.

doi:10.6342/NTU202304164

Decker, C., & Baade, A. (2016). Consumer perceptions of co-branding alliances: Organizational dissimilarity signals and brand fit. *Journal of brand management*, *23*, 648-665.

Freberg, K., Graham, K., McGaughey, K., & Freberg, L. A. (2011). Who are the social media influencers? A study of public perceptions of personality. *Public relations review*, 37(1), 90-92.

Han, H., Wang, W.-Y., & Mao, B.-H. (2005). Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. *Advances in Intelligent Computing, ICIC 2005* (pp. 878-887). Springer.

He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *Proceedings of 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)* (pp. 1322-1328).

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 770-778).

Heskett, J. L. (1976). *Marketing*. Macmillan.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, *9*(8), 1735-1780.

Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of International conference on machine learning, ICML 2015* (pp. 448-456).

Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing*, *57*(1), 1-22.

doi:10.6342/NTU202304164

Keller, K. L. (2013). *Strategic Brand Management: Building, Measuring, and Managing Brand Equity*. Pearson.

Keller, K. L. (2016). Reflections on customer-based brand equity: perspectives, progress, and priorities. *AMS review*, 6, 1-16.

Kerin, R. A., Harvey, M. G., & Rothe, J. T. (1978). Cannibalism and new product development. *Business Horizons*, 21(5), 25-31.

Koch, C., Lode, M., Stohr, D., Rizk, A., & Steinmetz, R. (2018). Collaborations on YouTube: From unsupervised detection to the impact on video and channel popularity. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(4), 1-23.

Koschmann, A. (2017). Brand alliances: Growing your pie or stealing your slice? examining field evidence using causal methods. *Journal of Marketing Development & Competitiveness*, 11(4).

Lassar, W., Mittal, B., & Sharma, A. (1995). Measuring customer-based brand equity. *Journal of consumer marketing*, 12(4), 11-19.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.

Lomax, W., Hammond, K., East, R., & Clemente, M. (1996). The measurement of cannibalization. *Marketing Intelligence & Planning*, 14, 20-28.

Ma, R., Gui, X., & Kou, Y. (2023). Multi-Platform Content Creation: The Configuration of Creator Ecology through Platform Prioritization, Content Synchronization, and Audience Management. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-19)

doi:10.6342/NTU202304164

Newmeyer, C. E., Venkatesh, R., & Chatterjee, R. (2014). Cobranding arrangements and partner selection: A conceptual framework and managerial guidelines. *Journal of the Academy of Marketing Science*, 42, 103-118.

Novelli, F. (2013). Measuring Sales Cannibalization in Information Technology Markets: Conceptual Foundations and Research Issues. In G. Herzwurm & T. Margaria, *Software Business. From Physical Products to Software Services and Solutions, ICSOB 2013* (pp. 31-42). Springer, Berlin, Heidelberg.

Paydas Turan, C. (2021). Success drivers of co-branding: A meta-analysis. *International Journal of Consumer Studies*, 45(4), 911-936.

Rao, A. R., Qu, L., & Ruekert, R. W. (1999). Signaling unobservable product quality through a brand ally. *Journal of marketing research*, 36(2), 258-268.

Ruekert, R. W., & Rao, A. (1994). Brand alliances as signals of product quality. *Sloan management review*, 36(1), 87-97.

Santurkar, S., Tsipras, D., Ilyas, A., & Madry, A. (2018). How does batch normalization help optimization? *Advances in neural information processing systems 31 (NeurIPS 2018)* (pp. 2488-2498).

Shocker, A. D., Srivastava, R. K., & Ruekert, R. W. (1994). Challenges and opportunities facing brand management: An introduction to the special issue. *Journal of marketing research*, 31(2), 149-158.

Simonin, B. L., & Ruth, J. A. (1998). Is a company known by the company it keeps? Assessing the spillover effects of brand alliances on consumer brand attitudes. *Journal of marketing research*, 35(1), 30-42.

doi:10.6342/NTU202304164

Van der Lans, R., Van den Bergh, B., & Dieleman, E. (2014). Partner selection in brand alliances: An empirical investigation of the drivers of brand fit. *Marketing science*, 33(4), 551-566.

Xiao, M., Wang, R., & Chan-Olmsted, S. (2018). Factors affecting YouTube influencer marketing credibility: a heuristic-systematic model. *Journal of media business studies*, 15(3), 188-213.

Zhou, R., Khemmarat, S., Gao, L., Wan, J., & Zhang, J. (2016). How YouTube videos are discovered and its impact on video views. *Multimedia Tools and Applications*, *75*, 6035-6058.

doi:10.6342/NTU202304164