mathematics

MDPI

Article

What Are the Public's Concerns About ChatGPT? A Novel Self-Supervised Neural Topic Model Tells You

Rui Wang [1,3], Xing Liu [1], Peng Ren [1], Shuyu Chang [1], Zhengxin Huang [2,*], Haiping Huang [1] and Guozi Sun [1]

[1] School of Computer Science, Nanjing University of Posts and Telecommunications, Nanjing 210023, China; hhp@njupt.edu.cn (H.H.)
[2] Department of Computer Science, Youjiang Medical University for Nationalities, Baise 533000, China
[3] Key Laboratory of New Generation Artificial Intelligence Technology and Its Interdisciplinary Applications (Southeast University), Ministry of Education, Nanjing 210023, China
* Correspondence: huangzhx26@mail2.sysu.edu.cn

Abstract: The recently released ChatGPT, an artificial intelligence conversational agent, has garnered significant attention in academia and real life. A multitude of early ChatGPT users have eagerly explored its capabilities and shared their opinions on social media, providing valuable feedback. Both user queries and social media posts have been instrumental in expressing public concerns regarding this advanced dialogue system. To comprehensively understand these public concerns, a novel Self-Supervised Neural Topic Model (SSTM), which formulates topic modeling as a representation learning procedure, is proposed in this paper. The proposed SSTM utilizes Dirichlet prior matching and three regularization terms for improved modeling performance. Extensive experiments on three publicly available text corpora (Twitter Posts, Subreddit and queries from ChatGPT users) demonstrate the effectiveness of the proposed approach in extracting higher-quality public concerns. Moreover, the SSTM performs competitively across all three datasets regarding topic diversity and coherence metrics. Based on the extracted topics, we could gain valuable insights into the public's concerns regarding technologies like ChatGPT, enabling us to formulate effective strategies to address these issues.

Keywords: text mining; information extraction; neural topic model; ChatGPT; social media analysis

MSC: 68T50

check for updates

Academic Editor: Jüri Majak

Received: 14 September 2024
Revised: 19 December 2024
Accepted: 29 December 2024
Published: 8 January 2025

Citation: Wang, R.; Liu, X.; Ren, P.; Chang, S.; Huang, Z.; Huang, H.; Sun, G. What Are the Public's Concerns About ChatGPT? A Novel Self-Supervised Neural Topic Model Tells You. *Mathematics* **2025**, *13*, 183. https://doi.org/10.3390/math13020183

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The emergence of ChatGPT [1] has significantly impacted dialogue generation with its remarkable performance. The public has widely embraced and actively participated in discussions about this chatbot [2,3] on social media like Twitter (http://twitter.com, accessed on 4 February 2024) and Reddit (http://reddit.com, accessed on 4 February 2024). For instance, Figure 1 shows the discussions regarding diverse concerns about ChatGPT which are posted by Twitter users. The distinct colors are marked to represent words for specific concerns (e.g., yellow means '*Education*' , green denotes '*Security*', and blue represents '*Technical Basis*'). Also, a vast number of users eagerly interact with this conversational agent and explore its capabilities. Compared to texts from other sources, the review posts and queries about ChatGPT encapsulate the primary concerns more effectively, such as the services it can offer and user experiences. Therefore, mining these silent concerns lying behind the texts is valuable for academia [4] and public daily life.

However, the massive volume of review posts and user-generated queries, the wide range of concerns on this open domain topic, and the noisy nature (e.g., casual expressions, grammar mistakes) [5] of the user-generated unlabeled text make it challenging to mine high-quality concerns. Considering the characteristics of this extraction task, topic models could provide a promising way to extract these public concerns in an unsupervised manner. Furthermore, to ensure the quality of the extracted concerns, we should consider both of the following aspects: (1). *Interpretability*: extracted concerns should be semantically coherent and agree with human judgment (mined representative words should be semantically coherent and have the ability to represent the meaning of specific concerns, such as the highlighted words in Figure 1). (2). *Diversity*: a wide range of concerns on ChatGPT should be reflected in mined results (all the discussed concerns, like '*Education*', '*Security*' and '*Technical Basis*' shown in Figure 1 should be discovered in the extracted results without neglecting any of them).

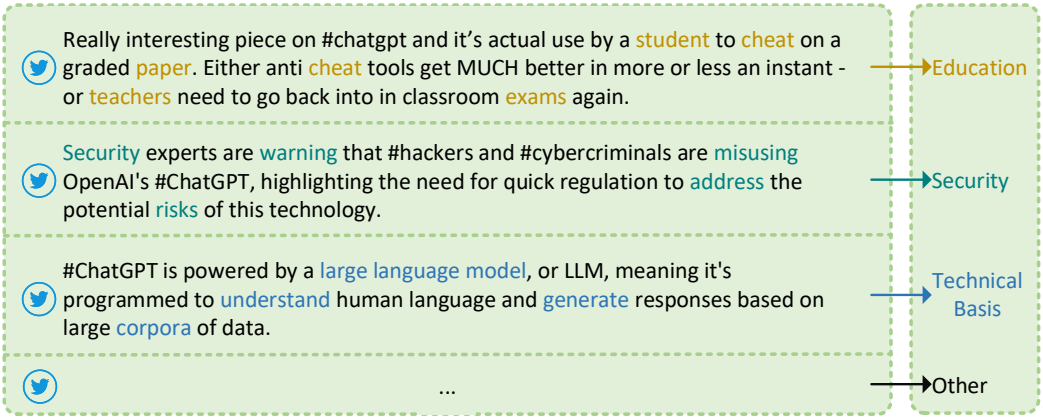


Figure 1. Three tweet examples of ChatGPT which convey different concerns; colored words are representative words of specific concerns, where yellow means '*Education*', green denotes '*Security*', and blue represents '*Technical Basis*'.

Neural topic modeling [6] addresses the inefficiency of the Gibbs sampler in traditional topic models [7], serving as a tool for unsupervised open information extraction [8] and offering a promising direction for mining public concern. Nevertheless, two streams of advanced models based on the Variational Auto-Encoder (VAE) [9] and Generative Adversarial Network (GAN) [10] could not fulfill both of our requirements. The improper prior of topic distribution used in VAE-based neural topic models [7,11–13] often sacrifices the interpretability, thereby diminshing the ability to represent specific focuses. Adversarial topic models [10,14] face topic collapse problems (fail in discovering diverse semantic patterns [7,15]), resulting in a reduction in concern diversity and the loss of crucial information.

Thus, to address the potential issues posed by the aforementioned approaches and mine high-quality public concerns about ChatGPT, the Self-Supervised Neural Topic Model (SSTM) is proposed in this paper. Compared to most existing models, topic modeling is formalized as representation learning to mitigate topic collapse [16]. Its principal idea is to build a one-way projection from the document-word distribution to the document-topic distribution, along with a text augmentation scheme, to capture the correlations between words and topics. Specifically, to ensure interpretability, we leverage an invariance regularizer to preserve topic information from augmented pairs and the maximum mean discrepancy to match the statistical distribution to the Dirichlet prior in the topic space. Additionally, a covariance regularizer is utilized to decorrelate dimensions of topic distributions and alleviate the generation of mixed topics. On the other hand, SSTM also

incorporates the variance regularizer to prevent topic collapse and ensure the diversity of concerns/topics.

The main contributions of our work are as follows:

- An end-to-end framework is proposed for mining public concerns about ChatGPT from social media posts and user queries, which is, to our best knowledge, the first attempt for this task.
- A novel Self-Supervised Neural Topic Model (SSTM) based on representation learning is proposed to mine high-quality concerns/topics.
- Experimental results on three publicly available datasets show that the SSTM could discover higher-quality concerns/topics than state-of-the-art approaches in terms of interpretability and diversity.

The practical significance of this research lies in designing the novel Self-Supervised Neural Topic Model for mining public concerns about the advanced AI-dialogue system ChatGPT. Additionally, it has the potential to generate high-quality public concerns when compared to existing state-of-the-art approaches. The remainder of the paper is organized as follows: Section 2 presents the related technical work, while Section 3 delves into the detailed architecture of the proposed SSTM. In Section 4, the experimental results and analysis, which include experimental setup (Section 4.1), interpretability and diversity evaluation (Section 4.2), hyper-parameter analysis (Section 4.3), human evaluation (Section 4.4), a case study (Section 4.5), data analysis (Section 4.6), and an ablation study (Section 4.7), are presented. Section 5 will discuss the results. At last, Section 6 will conclude this paper and provide future directions.

2. Related Work

Our work is closely related to two lines of researches, which are self-supervised learning and neural topic modeling.

2.1. Self-Supervised Learning

Self-supervised learning (SSL) is a family of techniques in machine learning that aims to learn high-quality representations from unlabeled data without relying on external supervision and human-labeled annotations. Generally, SSL designs pretext tasks [17] based on the inherent structure and patterns in the datasets to help models to capture high-quality representations [18].

SSL has made significant progress and is widely explored in the computer vision community due to its effectiveness in generating high-quality vision representations. Zbontar et al. [16] incorporated a redundancy reduction mechanism into SSL and proposed the Barlow Twins for vision representation learning. Bijoy et al. [19] developed an automated image cleaner in a few-shot learning scenario based on SSL. Fini et al. [20] incorporated the self-supervised objective into semi-supervised clustering to enforce the cluster centroids as class prototypes. Furthermore, Bardes et al. [21] proposed the VICRegL, which learns good global and local features simultaneously and yields an excellent performance on image segmentation.

Likewise, in the natural language processing (NLP) community, SSL has also been explored and has exhibited significant promise. Yan et al. [22] proposed the Contrastive framework for self-supervised SEntence Representation Transfer (ConSERT) for producing high-quality sentence representations. Zhu et al. [23] utilized the SSL for multi-modal sentiment analysis and image–text matching. Mu et al. [24] proposed an SSL-based approach for detecting synonym words from short texts. Luo et al. [25] used the GraphACL scheme to improve the performance of self-supervised graph representation learning. Yi et al. [26] propose a novel self-supervised deep graph clustering method (R^2FGC) to explore the

full semantic information of the graph-structured data. In addition, SSL has also made significant progress in other research on graph representation learning [27,28]. The latest advancements in self-supervised learning have significantly propelled the development of large language models (LLMs) [29,30], enabling them to learn from vast amounts of unlabeled data and improve their performance on various natural language processing tasks. However, no attempts have explored topic modeling using this advanced representation learning technique.

2.2. Neural Topic Modeling

Recently, Neural Topic Models (NTMs) have gained widespread attention as an efficient tool for mining high-quality topics from massive corpora. While traditional topic models like Latent Dirichlet Allocation (LDA) [31] face certain tricky issues, such as computational complexity and lack of scalability, NTMs offer a promising alternative.

The Variational Auto-Encoder (VAE) [9] has emerged as the mainstream framework for neural topic modeling. Inspired by the VAE, Miao et al. [32] proposed the Neural Variational Document Model (NVDM), which models latent topics with multivariate Gaussian priors and utilizes a decoder to reconstruct documents for text modeling. Likewise, Srivastava et al. [33] employed the Logistic-Normal distribution as the prior of topic space to mimic Dirichlet and proposed the ProdLDA based on the VAE. Along this line, Dieng et al. [12] modeled topics in the embedding space and proposed the Embedded Neural Topic Model, which could incorporate semantic relatedness in word embeddings into the modeling process. However, such VAE-based approaches have limitations in mining interpretable topics as they often use Gaussian and Logistic-Normal distributions as the prior distribution in latent topic space. These distributions could not capture multi-modality among texts and are unsuitable for text modeling.

On the other hand, Generative Adversarial Networks (GANs) have also been explored for neural topic modeling. Wang et al. [14] introduced adversarial learning to the topic model and proposed the Adversarial-neural Topic Model (ATM), which employs a generator to produce document-word distributions by sampling from a Dirichlet distribution, and a discriminator is utilized to differentiate real and generated documents. However, the ATM has limitations in obtaining the topic distribution for unseen documents. Building upon this, Wang et al. [10] proposed the Bidirectional Adversarial neural Topic (BAT) model. It incorporates an encoder to establish a bidirectional mapping between topic-word distributions and document-word distributions. Nevertheless, these adversarial topic models also have the following drawbacks: (1). GAN-based approaches often face topic collapse, which may result in a low level of topic diversity. (2). The alternate optimization procedure of adversarial training makes these models often encounter training instability problems.

Though neural topic modeling and self-supervised learning have been explored and made some progress, seldom have attempts been made to tackle the topic modeling problem using self-supervised learning. Thus, the proposed Self-Supervised Neural Topic Model (SSTM) differs from existing models in the following facets: (1). Unlike VAE-based approaches, which utilize improper prior distribution (such as Gaussian and Logistic-Normal distributions) in latent topic space, the SSTM models topics with the Dirichlet prior to ensure topic interpretibility. (2). Different from adversarial-based approaches that often face mode collapse and notoriously unstable training, our proposed SSTM formulates topic modeling as textual representation learning to mitigate topic collapse. (3). The SSTM incorporates variance and covariance regularizers to enhance topic diversity and alleviate the mixed topics by decorrelating the different dimensions in topic space.

3. Research Methodology

The proposed Self-Supervised Neural Topic Model (SSTM), shown in Figure 2, contains three main components: (1). The text augmentation and representation module (left), which defines the text augmentation and representation scheme. Given a batch of documents X, for each document $\vec{x} \in X$, it outputs document-word distributions of augmented texts $\vec{x}_a$ and $\vec{x}_b$; (2). The topic inference network I (middle) utilizes document-word distributions $\vec{x}_a$ and $\vec{x}_b$ as input and transforms them into document-topic distributions $\vec{\theta}_a$ and $\vec{\theta}_b$; (3). The prior matching module (right) aims to push the statistical distribution of document-topic distributions to the Dirichlet prior and calculate the representation learning objectives. The functionalities of these components will be introduced in the following subsections. Additionally, for the clarity of illustration, we list the symbols used in this paper and corresponding explanations in Table 1.

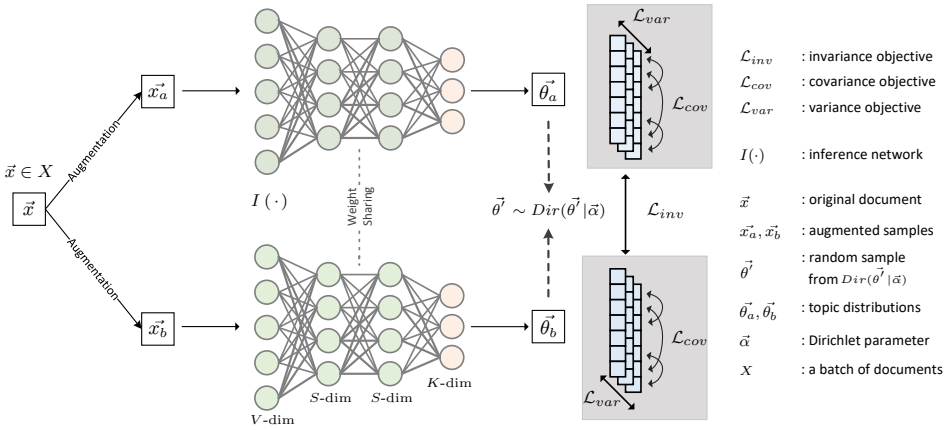


Figure 2. The overall framework of proposed Self-Supervised Neural Topic Model.

Table 1. The notations used in this paper and their descriptions.

Symbol	Description
	Data Representation and Distribution
X	a batch of documents
$\vec{x}$	a document represented with tf-idf
$\vec{x_a}, \vec{x_b}$	augmented documents of $\vec{x}$
$\vec{\theta''} \in \{\vec{\theta_a}, \vec{\theta_b}\}$	the inferred document-topic distributions for augmented documents
Θ	a set of inferred document-topic distributions
Θ'	a set of random samples drawn from the Dirichlet prior
$\vec{\alpha}$	the hyper-parameter of the Dirichlet prior distribution
$\mathbb{P}_{\Theta}$	the statistical distribution of inferred document-word distributions
$\mathbb{P}_{\Theta'}$	the Dirichlet prior distribution
	Model Parameter
I	the inference network
K, V, S	topic number, vocabulary size and the number of units in semantic representation layer
W_1, W_2, W_t	the weight matrices of different layers in I
$\vec{b_1}, \vec{b_2}, \vec{b_t}$	the bias terms in I
$\vec{h_1}, \vec{h_2}, \vec{h_t}$	hidden states of different layers in I
$\vec{o}_1, \vec{o}_2$	activated signals of hidden layers in I
$\mathrm{SN}(\cdot), \mathrm{BN}(\cdot)$	the spectrum normalization and the batch normalization
$\mathrm{Hardswish}(\cdot)$	the Hardswish activation function
$\kappa(\cdot, \cdot)$	the kernel function
$\mathcal{L}_{inv}$	the invariance objective
$\mathcal{L}_{var}$	the variance objective
$\mathcal{L}_{cov}$	the covariance objective
$\mathcal{L}_R$	the representation learning objective
$\mathcal{L}_M$	the prior matching objective

Table 1. *Cont.*

Symbol	Description
	Parameter of Training Procedure
E, M, B	number of epoch, batch size and number of batch
W	V-dimensional diagonal matrix
$\beta, \lambda, \mu, \eta$	the coefficient of regularizers
Φ	correlation matrix between words and topics
$\Phi^{'}$	topic-word distributions matrix
p	probability used in text augmentation

3.1. Text Augmentation and Representation

Text augmentation plays an essential role in self-supervised representation learning, and its key idea is keeping the semantic consistency between augmented text and original text.

Thus, for each document with tf-idf representation $\vec{x}$, we conduct various random perturbations to generate augmented samples that share the same semantic meaning with $\vec{x}$. To carry out the augmentation, we first select L words (set to 5 in our experiments) that appear in $\vec{x}$ with the lowest weights. Then, we conduct the following perturbations for each selected word:

1. Increase weight by 10% with p% probability (p is set to 15 in the experiment);
2. Decrease weight by 10% with p% probability (p is set to 15 in the experiment);
3. Set weight to zero with p% probability (p is set to 15 in the experiment).

Then, we represent the augmented texts with document-word distributions $\vec{x}_a$ and $\vec{x}_b$, which can be obtained by dividing the summation of weights over vocabulary.

3.2. Topic Inference Network

The topic inference network I aims to capture the correlations between words and topics by building the projection from the document-word distribution to the document-topic distribution. It consists of one V-dimensional document-word distribution layer, two S-dimensional semantic representation layers, and one K-dimensional document-topic distribution layer.

In more detail, for each document-word distribution $\vec{x}^{'} \in \{\vec{x}_a, \vec{x}_b\}$, inference network I transforms it into the S-dimensional semantic space at first with the transformation defined as:

$$\vec{h}_1 = \text{SN}(W_1 \vec{x}^{'} + \vec{b}_1) \tag{1}$$

$$\vec{o}_1 = \text{Hardswish}(\vec{h}_1) \tag{2}$$

$$\vec{h}_2 = \text{SN}(W_2 \vec{o}_1 + \vec{b}_2) \tag{3}$$

$$\vec{o}_2 = \text{Hardswish}(\vec{h}_2) \tag{4}$$

where $W_1 \in \mathbb{R}^{S \times V}$ and $W_2 \in \mathbb{R}^{S \times S}$ are weight matrices of semantic layers and $\vec{b}_1$ and $\vec{b}_2$ are bias terms. $\vec{h}_1$, $\vec{h}_2$, $\vec{o}_1$, and $\vec{o}_2$ are hidden states and activation signals of the two semantic layers. Moreover, $\text{SN}(\cdot)$ denotes spectrum normalization [34], and $\text{Hardswish}(\cdot)$ is the Hardswish activation function [35]. Then, the topic inference network will project the output signal $\vec{o}_2$ into a K-dimensional document-topic distribution via the formulation below:

$$\vec{h}_t = \text{BN}(W_t\vec{o}_2 + \vec{b}_t) \tag{5}$$

$$\vec{\theta''} = \text{softmax}(\vec{h}_t) \tag{6}$$

where $W_t \in \mathbb{R}^{K\times S}$ and $\vec{b}_t$ are weight matrix and bias terms, $\vec{h}_t$ means the hidden states of the document-topic distribution layer, and $\vec{\theta''} \in \{\vec{\theta}_a, \vec{\theta}_b\}$ denotes the inferred topic distribution.

3.3. Prior Matching in Topic Space

As Wallach et al. [36] argued that multiple peaks of the Dirichlet density are suitable for mining multi-modality among texts, we model topics with the Dirichlet prior in latent topic space.

To match topic distributions to the pre-defined Dirichlet prior, we utilize the Maximum Mean Discrepancy (MMD) [37], which is a widely used statistical distribution matching tool. MMD measures the distance between probability densities [38] by mapping their mean embeddings into a Reproducing Kernel Hilbert Space (RKHS) [39]. Concretely, given a set of inferred topic distributions $\Theta=\{\vec{\theta}_a^1, \ldots, \vec{\theta}_a^M, \vec{\theta}_b^1, \ldots, \vec{\theta}_b^M\}$ sampled from their statistical distribution $\mathbb{P}_\Theta$ and a set of random samples $\Theta'=\{\vec{\theta}'_1, \vec{\theta}'_2, \ldots, \vec{\theta}'_{2M}\}$ drawn from the Dirichlet prior distribution $\mathbb{P}_{\Theta'}$, the distance between $\mathbb{P}_\Theta$ and $\mathbb{P}_{\Theta'}$ can be calculated with:

$$\mathcal{D}(\mathbb{P}_\Theta, \mathbb{P}_{\Theta'}) = \frac{1}{2M(2M-1)}\sum_{i\neq j}(\kappa(\vec{\theta}_i, \vec{\theta}_j) + \kappa(\vec{\theta}'_i, \vec{\theta}'_j)) - \frac{1}{2M^2}\sum_{i=1}^{2M}\sum_{j=1}^{2M}\kappa(\vec{\theta}_i, \vec{\theta}'_j) \tag{7}$$

where $\kappa(\cdot,\cdot)$ denotes the kernel function. In this paper, we use the diffusion kernel [40] and M represents the batch size.

3.4. Training Objective

The principal idea of the proposed SSTM could be summarized as:

- Utilizing the topic inference network to build the projection from the document-word distribution to the document-topic distribution and learn the high-quality topic distribution;
- Matching the statistical distribution $\mathbb{P}_\Theta$ of inferred document-topic distributions to the Dirichlet prior $\mathbb{P}_{\Theta'}$ with MMD to capture multiple patterns.

Thus, to make sure that the SSTM can mine high-quality concerns considering both interpretability and diversity, its training objective is designed as:

$$\mathcal{L} = \mathcal{L}_R + \beta\mathcal{L}_M \tag{8}$$

where $\mathcal{L}_M$ means the prior matching objective, which is calculated with the MMD distance defined in Equation (7), and β denotes the coefficient of the prior matching regularizer.

On the other hand, to learn a high-quality projection and provide informative topic distributions, the representation learning objective $\mathcal{L}_R$ should consider the requirements below:

- Invariance: The inference network should map similar document-word distributions ($\vec{x}_a$ and $\vec{x}_b$) to similar document-topic distributions ($\vec{\theta}_a$ and $\vec{\theta}_b$).
- Variance: Topic collapse should be prevented as it will map different document-word distributions to the same document-topic distribution and deteriorate topic diversity.
- Covariance: Different dimensions of inferred document-topic distributions should be decorrelated, and each dimension should capture an independent semantic meaning lying behind texts.

Given two batches of inferred topic distributions $\Theta_a = \{\vec{\theta}_a^1, \ldots, \vec{\theta}_a^M\}$ and $\Theta_b = \{\vec{\theta}_b^1, \ldots, \vec{\theta}_b^M\}$, we should first ensure that similar document-word distributions have similar document-topic distributions (invariance) by calculating the averaged squared Euclidean distance

between topic distribution pairs. By minimizing the Euclidean distance between topic distributions of paired augmented documents, we could ensure that semantically similar documents have similar document-topic distributions and enhance the topic interpretability. The invariance objective function $\mathcal{L}_{inv}$ is defined as:

$$\mathcal{L}_{inv} = \frac{1}{M}\sum_{i=1}^{M} ||\vec{\theta}_a^i - \vec{\theta}_b^i||_2^2 \tag{9}$$

where M denotes the batch size.

Then, to prevent topic collapse, we incorporate a variance regularization term, and the variance objective function of Θ_a is formed as:

$$\mathcal{L}_{var}(\Theta_a) = \frac{1}{K}\sum_{k=1}^{K} \max(0, 1 - \varsigma(\Theta_a^k, \epsilon)) \tag{10}$$

where $\varsigma(\vec{z}, \epsilon) = \sqrt{\mathrm{Var}(\vec{z}) + \epsilon}$, Θ_a^k represents the vector composed of each value at the k-th dimension in all topic distributions in Θ_a, and ϵ is a small scalar which is set to 1×10^{-4} for numerical stability. This regularizer enforces a variance of 1 for each dimension within the current batch, helping to prevent all inputs from collapsing onto a single vector. Also, this term ensures the variance of each dimension of topic distribution and improves the topic diversity. The variance objective function of Θ_a and Θ_b can be summarized as $\mathcal{L}_{var}$ with the formulation:

$$\mathcal{L}_{var} = \frac{1}{2}(\mathcal{L}_{var}(\Theta_a) + \mathcal{L}_{var}(\Theta_b)) \tag{11}$$

where $\mathcal{L}_{var}(\Theta_b)$ follows the formula in Equation (10).

Moreover, to disentangle the semantic meaning of topics, the covariance objective function is designed based on the covariance matrix which is formulated as:

$$C(\Theta_a) = \frac{1}{M-1}\sum_{i=1}^{M}(\vec{\theta}_a^i - \boldsymbol{\theta}_a)(\vec{\theta}_a^i - \boldsymbol{\theta}_a)^\top \tag{12}$$

where $\boldsymbol{\theta}_a = \frac{1}{M}\sum_{i=1}^{M}\vec{\theta}_a^i$. Aiming at decorrelating the dimensions of document-topic distributions, we follow VICREG [41] and define a covariance objective function $\mathcal{L}_{cov}$ with:

$$\mathcal{L}_{cov} = \frac{1}{2K}(\sum_{i\neq j}[C(\Theta_a)]_{i,j}^2 + \sum_{i\neq j}[C(\Theta_b)]_{i,j}^2) \tag{13}$$

where $C(\Theta_b)$ follows the calculation in Equation (12). The covariance objective helps SSTM disentangle the dimensions of document-topic distributions and improve topic interpretability.

With the definition of invariance objective function $\mathcal{L}_{inv}$ (Equation (9)), variance objective function $\mathcal{L}_{var}$ (Equation (11)), and covariance objective function $\mathcal{L}_{cov}$ (Equation (13)), the representation learning objective function $\mathcal{L}_R$ can be formulated with:

$$\mathcal{L}_R = \lambda\mathcal{L}_{var} + \mu\mathcal{L}_{cov} + \eta\mathcal{L}_{inv} \tag{14}$$

where λ, μ and η are coefficients of different regularization terms. The training procedure of the proposed SSTM is shown in Algorithm 1. For detailed configurations, refer to Section 4.1.4.

Algorithm 1: Self-Supervised neural Topic Model

Input: Dataset $D = \{\vec{x_1}, \ldots, \vec{x_N}\}$, topic number K, number of epoch E, batch size M, number of batch B in the epoch, coefficient parameters $[\beta, \lambda, \mu, \eta]$.
Output: The learned inference network I.
Initialize the inference network I and shuffled dataset D.
for *each epoch in E* **do**
  for *each batch in B* **do**
    Draw a batch M of documents X from D
    Augment X and obtain $X^{'} = X_a \cup X_b = \{\vec{x}_a^1, \ldots, \vec{x}_a^M, \vec{x}_b^1, \ldots, \vec{x}_b^M\}$
    Obtain topic distributions $\Theta = \Theta_a \cup \Theta_b$ with $\Theta_a = I(X_a)$ and $\Theta_b = I(X_b)$
    Draw two batches $2M$ of random samples $\Theta' \sim Dir(\vec{\theta}'|\vec{\alpha})$
    Compute training objectives:
      Prior Matching objective $\mathcal{L}_M$ by Equation (7)
      Invariance objective $\mathcal{L}_{inv}$ by Equation (9)
      Variance objective $\mathcal{L}_{var}$ by Equations (10) and (11)
      Covariance objective $\mathcal{L}_{cov}$ by Equations (12) and (13)
    Compute the representation learning objective $\mathcal{L}_R$ with Equation (14)
    Compute the overall training objective: $\mathcal{L} = \mathcal{L}_R + \beta\mathcal{L}_M$
    Update I and minimize $\mathcal{L}$ with gradient back-propagation

3.5. Topic Generation

After model training, the inference network I builds the projection from the document-word distribution layer to the document-topic distribution layer, which could be utilized for topic mining.

By feeding into the diagonal matrix $W \in \mathbb{R}^{V \times V}$ (each row represents the one-hot encoding of a word in the vocabulary), we can obtain correlations between topics and words via the transformation below:

$$\Phi = I(W) \tag{15}$$

where $\Phi \in \mathbb{R}^{V \times K}$ is the correlation matrix between words and topics. By performing column-wise normalization and ranking, we can obtain the topic-word distribution matrix $\Phi^{'}$ (k-th column is the topic-word distribution of the k-th topic), which can be used for topic extraction.

4. Experiments

In this section, we will first describe the experimental setup. Then, interpretability and diversity evaluation, a case study, hyper-parameter analysis, and an ablation study will be presented.

4.1. Experimental Setup

In this subsection, we will first provide descriptions of datasets and baselines. Following this, evaluation metrics and implementation details will be introduced.

4.1.1. Datasets

To evaluate the performance of the proposed model, we collect and process three publicly available datasets, which are the Twitter Posts (https://www.kaggle.com/datasets/manishabhatt22/tweets-onchatgpt-chatgpt, accessed on 8 February 2024), Subreddit (https://www.kaggle.com/datasets/armitaraz/chatgpt-reddit, accessed on 8 February 2024), and

User Query (https://www.kaggle.com/datasets/noahpersaud/89k-chatgpt-conversations, accessed on 8 February 2024) datasets. In more detail, the Twitter Posts dataset contains 58 K tweets with the hashtag '*#ChatGPT*' posted on Twitter from 30 November 2022 to 24 February 2023. The Subreddit dataset consists of nearly 18 K reviews regarding ChatGPT posted on the Reddit website. Both datasets encompass diverse user attitudes and opinions on this AI-dialogue system and are suitable for public opinion mining. When collecting these data, we define the scope based on the communities of ChatGPT, without considering factors such as the geographical location and age range of the commenters. On the other hand, the User Query dataset contains all the available conversations from chatlogs.net between users and ChatGPT until 20 April 2023. This dataset is collected from chatlogs.net using a custom web scraper (https://github.com/P1ayer-1/chatlogs.net-scraper, accessed on 20 February 2024) and its content contains user-concerned aspects and expectations of this powerful tool. In our experiment, we conduct an anonymization process, such as removing and replacing personally identifiable information, to protect personal privacy. Additionally, to address data noise, we conduct data preprocessing before performing experiments. To remove tweet-specific noise (such as hashtags, URLs, numbers, etc.), we first discard the non-English tweets with *fastlid* (https://pypi.org/project/fastlid/, accessed on 25 February 2024). Then, we employ the *preprocessor* (https://github.com/s/preprocessor, accessed on 25 February 2024) and *twitter_preprocessor* (https://github.com/vasisouv/tweets-preprocessor, accessed on 25 February 2024) libraries for cleaning. We also employ *enchant* (https://github.com/AbiWord/enchant, accessed on 25 February 2024) and *spacy* (https://github.com/explosion/spaCy, accessed on 25 February 2024) for spell-checking and lemmatization. The stopwords and low-frequency words are removed. Table 2 presents the statistics of the processed datasets. Likewise, we also provide the word-cloud visualizations of three datasets in Figure 3.

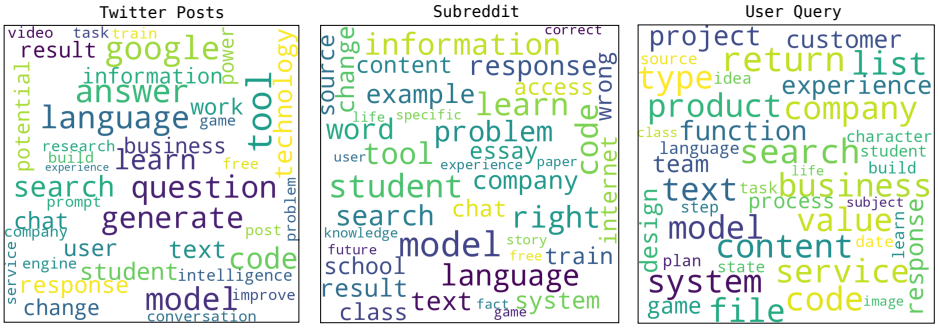


Figure 3. Word-cloud visualizations of three datasets.

Table 2. Statistics of the processed datasets.

Dataset	#Documents	Avg Length	Vocab Size
Twitter Posts	57,724	14.30	6437
Subreddit	17,438	25.31	4221
User Query	18,258	66.56	5055

4.1.2. Baselines

To validate the effectiveness of the SSTM, we choose the following methods as baselines:

- **LDA** [31], a probabilistic graphical model which views document is generated by a mixture of topics, and we employ the *Mallet* (https://mimno.github.io/Mallet/, accessed on 15 March 2024) implementation with suggested configurations.

- **NVDM** [32], a VAE-based neural topic model which models topics with the Gaussian prior in latent topic space, official implementation is used in the experiments (https://github.com/ysmiao/nvdm, accessed on 15 March 2024).
- **GSM** [42], a variant of the NVDM, utilizes a Gaussian distribution followed by a softmax transformation to construct the topic distribution; the official implementation (https://github.com/linkstrife/NVDM-GSM, accessed on 15 March 2024) is used.
- **ProdLDA** [33], a neural topic model which uses the Logistic-Normal distribution as prior. The official implementation (https://github.com/akashgit/autoencoding_vi_for_topic_models, accessed on 16 March 2024) is utilized.
- **WLDA** [43], a neural topic model built upon the Wasserstein Auto-Encoder, utilizing the Dirichlet distribution as the prior of the topic distribution; the official implementation (https://github.com/awslabs/w-lda, accessed on 17 March 2024) is used.
- **ATM** [14], a neural topic model employing adversarial training; we implement the model according to the recommended configurations.
- **BAT** [10], an extension of ATM which incorporates an encoder network for topic inference; we implement it with the suggestions in the original paper.
- **ETM** [12], a VAE-based approach that leverages topic embedding and word embeddings to capture semantic relationships; the official implementation (https://github.com/lffloyd/embedded-topic-model, accessed on 17 March 2024) is used.
- **CNTM** [7], a neural topic modeling approach based on contrastive learning in latent topic space, and the official implementation is utilized (https://github.com/SZU-AdvTech-2022/160-Contrastive-Learning-for-Neural-Topic-Model, accessed on 20 March 2024).
- **BERTopic** [13], a topic extraction tool based on the pre-trained language model (BERT); the released source code is used (https://github.com/MaartenGr/BERTopic, accessed on 22 March 2024) in our experiments.
- **CTMNeg** [11], a topic mining tool that incorporates contrastive learning and negative sampling to effectively capture the underlying structure and semantics; the official implementation (https://github.com/adhyasuman/ctmneg, accessed on 22 March 2024) is used in this paper.

4.1.3. Evaluation Metrics

As interpretability and diversity are two crucial aspects of public concern extraction, we consider two types of metrics for performance evaluation. To evaluate the interpretability of extracted concerns, topic coherence [44] (C_P, NPMI, UCI) is employed as [45] argued that coherence metric agrees with human judgment. All these values are computed with Palmetto library (https://github.com/dice-group/Palmetto, accessed on 12 April 2024). For evaluating concern diversity, we utilize the Unique Term Rate (UTR), which is calculated with $\frac{|UT|}{K \times 10}$, where $|UT|$ denotes the number of unique words in the extracted topic words. Note that higher values refer to a better quality of concern.

4.1.4. Implementation Details

We adopt the training settings with 50 epochs and a batch size of 32 for different topic numbers $K \in [10, 20, 30, 40]$. The Dirichlet hyper-parameter $\vec{\alpha}$ is set to $0.\vec{0}1$, and each semantic representation layer contains 400 units. Furthermore, the coefficient weight μ, λ and η are set to 1, $K/2$, and K, respectively. Our proposed SSTM is optimized using the ADAM [46] optimizer with a learning rate lr of 2×10^{-5}. In our experiments, each topic is represented with a list of the top 10 words for interpretability and diversity evaluation.

Also, our source code relies on libraries: *PyTorch* (https://pytorch.org/, accessed on 8 March 2024), *SciPy* (https://scipy.org/, accessed on 8 March 2024), *NumPy* (https://numpy.

org/, accessed on 10 March 2024), *torch_two_sample* (https://torch-two-sample.readthedocs.io/en/latest/, accessed on 10 March 2024), and *scikit-learn* (https://scikit-learn.org/stable/, accessed on 10 March 2024). All of our experiments are carried out on a Ubuntu machine equipped with a single 3090Ti GPU (CUDA version 11.7).

4.2. Interpretability and Diversity Evaluation

We have conducted a comprehensive set of experiments to assess the performance of the proposed SSTM. To assess the interpretability of the extracted concerns, we first compare the averaged topic coherence and provide a global view of comparisons. The results are listed in Table 3. Each value is obtained by averaging the coherence values on 10, 20, 30, and 40 topic number settings. To clarify the comparison, the optimal result is marked in bold, and the values marked with underline perform the second best. The statistics reveal that the SSTM outperforms all the baselines on the Twitter Posts and User Query datasets. For the Subreddit dataset, the SSTM performs the best in terms of the NPMI and UCI metrics. However, it slightly falls behind CNTM on the C_P metric.

Table 3. Averaged topic coherence with four topic settings [10, 20, 30, 40] and the relative improvement rate. (The relative improvement is computed as $\frac{R_b - R_s}{R_b - R_w}$%. Define R_b as the best result, R_s as the second-best result, and R_w as the worst result).

Model	Twitter Posts				Subreddit				User Query			
	C_P	NPMI	UCI	Rank	C_P	NPMI	UCI	Rank	C_P	NPMI	UCI	Rank
LDA	0.1796	0.0274	0.0364	2	0.1816	0.0334	0.1967	2	0.2103	0.0395	0.3096	3
NVDM	0.0868	0.0075	−0.0357	8	0.0943	0.0184	0.1373	6	0.1166	0.0107	−0.0208	8
WLDA	0.1607	0.0131	−0.4369	10	0.1892	0.0151	−0.5806	8	0.2008	0.0174	−0.4538	9
GSM	0.0210	−0.0209	−0.6608	12	0.0224	0.0121	−0.5007	11	0.0495	−0.0153	−0.5682	11
ProdLDA	0.1636	0.0196	−0.2222	6	0.1711	0.0172	−0.2856	7	0.1350	0.0041	−0.6050	10
ATM	0.1059	0.0050	−0.2510	11	0.0992	0.0110	−0.1073	12	0.1629	0.0189	−0.0815	7
BAT	0.1738	0.0223	−0.0791	4	0.1847	0.0285	0.0559	4	0.2123	0.0246	−0.2201	6
ETM	0.1439	0.0188	−0.2885	9	0.1434	0.0234	−0.2135	5	0.2247	0.0397	0.0985	2
CNTM	0.1589	0.0176	−0.2257	7	**0.2519**	0.0328	−0.2509	3	0.2834	0.0394	−0.0779	4
BERTopic	0.2177	0.0201	−0.6048	5	0.1117	0.0144	−1.0330	10	0.2777	0.0349	−0.2977	5
CTMNeg	0.1927	0.0285	−0.0363	3	0.0387	0.0369	−1.4599	9	0.0745	−0.0271	−1.6668	12
SSTM	**0.2196** ↑ 1%	**0.0401** ↑ 19%	**0.1555** ↑ 14.6%	1	0.2321 ↓ −8.6%	**0.0430** ↑ 30%	**0.2401** ↑ 2.6%	1	**0.2983** ↑ 6%	**0.0581** ↑ 21.6%	**0.4084** ↑ 4.8%	1

We can observe that approaches like CNTM and BERTopic perform well on the User Query dataset while performing worse on others. This may be attributed to the fact that they could not deal with the data sparsity of short-length social media posts. As queries often have more sufficient word-co-occurrence information compared with the Twitter and Reddit posts, CNTM and Bertopic perform well on the User Query dataset. With these observations, we could conclude that approaches are more prone to mining high quality topics/concerns from long text corpora.

In Figure 4, we also present the comparison of interpretability (topic coherence) vs. different topic numbers on Twitter Posts, Subreddit, and User Query datasets. For the Twitter Posts dataset, we can observe that the SSTM surpasses almost all the baselines except for the 10-topic number setting on UCI and the 40-topic number setting on C_P metrics. Similarly, it achieves optimal results across all the topic settings on three metrics for the User Query dataset, except for 40 topics on the UCI metric. Moreover, the SSTM outperforms most baselines on the NPMI and UCI metrics across different topics, but it obtains slightly worse results on the C_P metric. To provide a more intuitive comparison of extracted concerns, Table 4 presents the concern-specific representative words, obtained by the SSTM and baseline approaches. These concerns may be related to '*education*' and '*marketing*'.

On the other hand, we also provide a comparison of concern diversity across various topic numbers in Table 5. The experimental results demonstrate that our proposed SSTM

surpasses the baselines in all topic number settings on the Twitter Posts dataset. It also performs the best except for 30 and 40 topic settings on Subreddit and User Query. Notably, GSM demonstrates favorable performance on both datasets. However, it sacrifices the performance for topic coherence metrics, as depicted in Figure 4.

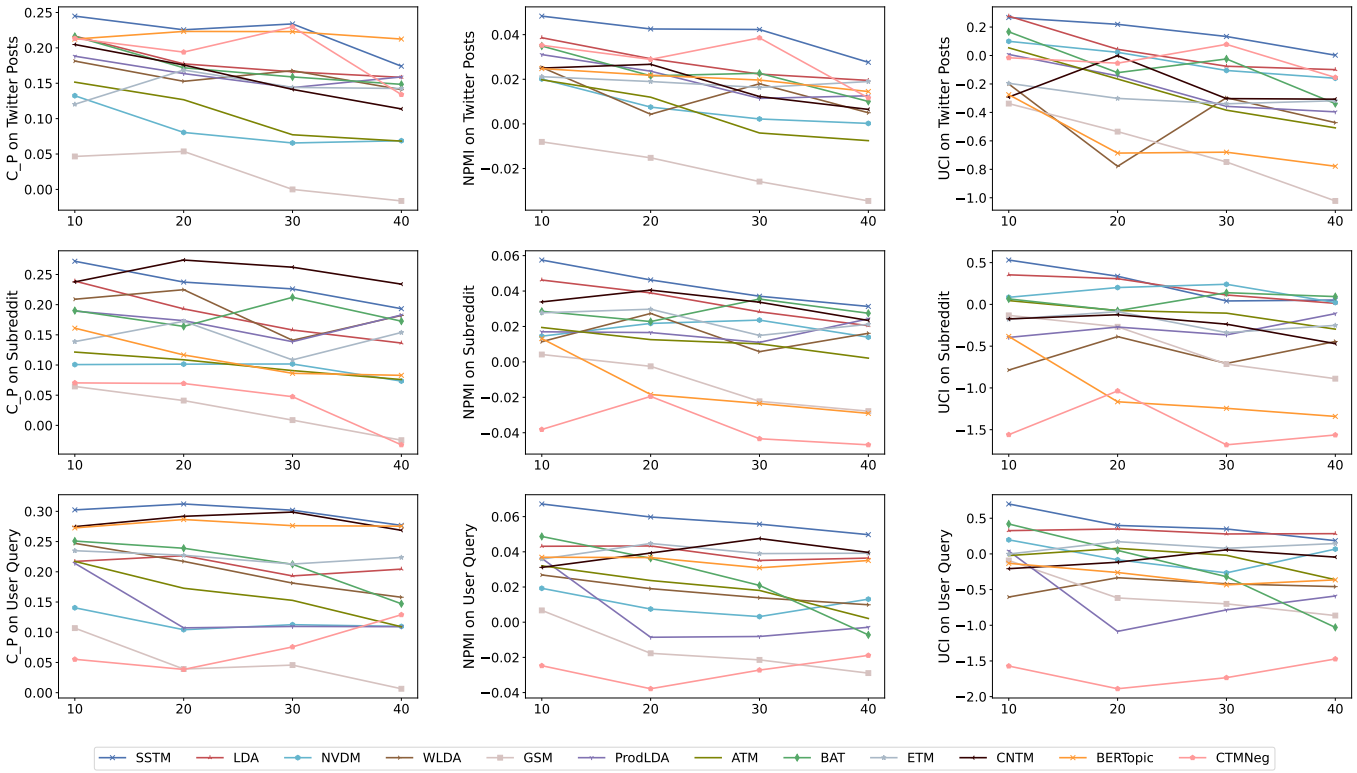


Figure 4. The comparison of interpretability (topic coherence) with different topic numbers on Twitter Posts, Subreddit, and User Query datasets.

Table 4. Concern (topic) examples extracted by SSTM and baselines from User Query dataset. Italics indicate out-of-topic word.

ATM		**BAT**		**CNTM**		**Bertopic**		**SSTM**	
Education	Marketing	Education	Marketing	Education	Marketing	Education	Marketing	Education	Marketing
student	product	student	market	student	bank	student	*team*	student	product
school	*month*	school	company	school	fund	*topic*	management	school	business
study	service	*experience*	*month*	learn	*rate*	*sentence*	business	research	company
university	company	skill	business	academic	banking	*question*	marketing	paper	team
teacher	platform	learn	investment	course	financial	*essay*	customer	essay	customer
learn	team	*team*	*water*	teacher	market	learn	product	course	market
science	business	research	cost	research	payment	*attribution*	development	learn	service
subject	marketing	*development*	*home*	education	investment	course	platform	study	marketing
course	digital	*project*	team	university	inflation	thesis	service	skill	sale
teach	client	*design*	*insurance*	classroom	invest	*language*	company	teacher	management

Generally, such improvements in topic coherence and diversity may be attributed to the incorporation of Dirichlet priors in latent space, invariance, and covariance regularizers. These comparisons also prove that the variance regularizer helps to address the topic collapse issue. Additionally, for the same metrics, approaches typically achieve better results on the User Query dataset, which may be caused by data characteristics such as longer text length and richer semantic information.

Table 5. The comparison of diversity (Unique Term Rate). The best result is marked in **bold**, and the second best is marked with an underline.

Model	Twitter Posts				Subreddit				User Query			
	10	20	30	40	10	20	30	40	10	20	30	40
LDA	0.960	0.900	0.913	0.887	0.920	0.975	0.960	0.887	0.960	0.960	0.973	0.948
NVDM	0.660	0.600	0.667	0.608	0.530	0.430	0.370	0.395	0.710	0.480	0.377	0.343
WLDA	0.900	0.745	0.550	0.405	0.910	0.615	0.580	0.520	0.970	0.800	0.603	0.530
GSM	**1.000**	**1.000**	0.970	0.963	**1.000**	0.980	**0.980**	**0.995**	**1.000**	**1.000**	**1.000**	**0.998**
ProdLDA	**1.000**	0.965	0.887	0.790	0.920	0.765	0.587	0.522	0.990	0.870	0.813	0.740
ATM	0.930	0.755	0.803	0.750	0.920	0.755	0.730	0.675	0.980	0.855	0.803	0.733
BAT	0.720	0.590	0.600	0.645	0.780	0.750	0.607	0.608	0.830	0.785	0.717	0.647
ETM	0.950	0.960	0.943	0.855	0.910	0.880	0.780	0.695	0.960	0.935	0.863	0.802
CNTM	0.990	0.530	0.417	0.355	0.950	0.910	0.870	0.785	**1.000**	0.890	0.787	0.703
BERTopic	0.830	0.900	0.847	0.910	0.840	0.815	0.807	0.767	0.900	0.915	0.917	0.902
CTMNeg	0.970	0.810	0.713	0.492	0.930	0.880	0.753	0.713	0.960	0.825	0.817	0.860
SSTM	**1.000**	**1.000**	**0.973**	**0.978**	**1.000**	**1.000**	0.967	0.922	**1.000**	**1.000**	0.990	0.958

4.3. Hyper-Parameter Analysis

To explore whether the performance of SSTM exhibits variations with changes in hyper-parameters (e.g., the learning rate lr, the Dirichlet prior $\vec{\alpha}$, the number of hidden units S, and the probability p in text augmentation), we conduct the hyper-parameter analysis in this subsection.

Specifically, considering practical issues such as experiment duration, we opted to conduct this experiment on the User Query dataset with the topic number set to 10. Likewise, C_P, UCI, and NPMI are utilized to measure the extracted concerns. Extensive experiments have demonstrated that the proposed SSTM performs stably when the parameters vary within the specified range, and experimental results are listed in Table 6. Furthermore, to present a more straightforward comparison, we visualize the results in Figure 5, and other parameters follow the default configurations. We will discuss the parameter analysis in more detail below.

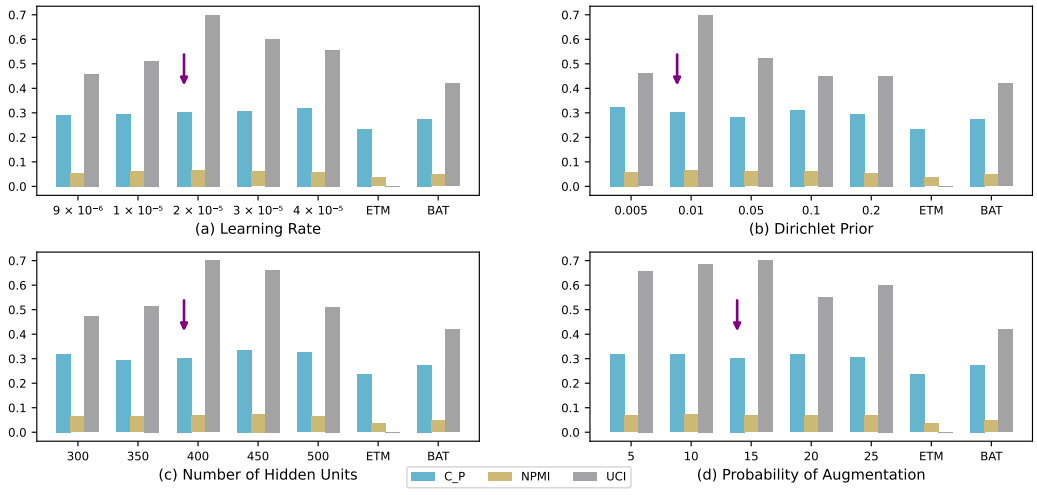


Figure 5. Performance analysis on various lr, $\vec{\alpha}$, S, and p. (We only retain ETM (second best for four topic number settings according to Table 3) and BAT (performs the second best with the topic number set to 10) for clarity).

Learning Rate (lr): To validate if the performance of the SSTM varies with different learning rates lr, we conduct parameter analysis experiments on lr, which is set to five values $\{9 \times 10^{-6}, 1 \times 10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}$ and $4 \times 10^{-5}\}$. The topic coherence results, listed in Table 6, show that the proposed SSTM outperforms baseline approaches with all

five settings of learning rate. Additionally, it performs the best when lr is set to 2×10^{-5}. Additionally, the comparison in Figure 5a proves that the SSTM is insensitive to the choice of learning rate parameter.

Dirichlet Prior ($\vec{\alpha}$): As the Dirichlet distribution plays a crucial role in mining multiple patterns for topic models, we also carry out the hyper-parameter $\vec{\alpha}$ of the Dirichlet density, and its values are set to {0.005, 0.01, 0.05, 0.1, 0.2}. Table 6 presents the results of the SSTM with various settings of $\vec{\alpha}$. We can observe that the SSTM performs stably across all five settings and it obtains the optimal coherence when $\vec{\alpha}$ is set to $\vec{0}.01$. The visualization comparison in Figure 5b reveals the same finding.

Number of Hidden Units (S): Likewise, to investigate the influence of the number of hidden units S on extraction performance, we employ five different settings {300, 350, 400, 450, 500} of S and conduct the corresponding analysis. As illustrated in Table 6, the SSTM with different numbers of hidden units outperforms the baseline approaches and performs the best when S is set to 450. Moreover, Figure 5c presents a more intuitive comparison.

Probability of Augmentation (p): To explore the impact of probability p utilized in text augmentation operation, we conducted experiments with five different settings $p \in \{5, 10, 15, 20, 25\}$. As shown by the statistics in Table 6, through variations in text augmentation probabilities, the proposed SSTM exhibits remarkable stability and consistently surpasses the competitive baseline, and it performs the best when p is set to 10. Likewise, the findings are also presented in Figure 5d.

Table 6. Topic coherence analysis of SSTM with various lr, $\vec{\alpha}$, S, and p. (We only retain ETM (second best for four topic number settings according to Table 3) and BAT (performs the second best with the topic number set to 10) for clarity).

Learning Rate lr							
Setting	9×10^{-6}	1×10^{-5}	2×10^{-5}	3×10^{-5}	4×10^{-5}	ETM	BAT
C_P	0.2920	0.2955	0.3024	0.3081	**0.3212**	0.2349	0.2509
NPMI	0.0549	0.0615	**0.0672**	0.0606	0.0600	0.0359	0.0487
UCI	0.4597	0.5124	**0.6998**	0.6009	0.5568	−0.0016	0.4192
Dirichlet Prior $\vec{\alpha}$							
Setting	0.005	0.01	0.05	0.1	0.2	ETM	BAT
C_P	**0.3236**	0.3024	0.2819	0.3105	0.2956	0.2349	0.2509
NPMI	0.0600	**0.0672**	0.0607	0.0604	0.0548	0.0359	0.0487
UCI	0.4603	**0.6998**	0.5240	0.4507	0.4491	−0.0016	0.4192
Hidden Layer S							
Setting	300	350	400	450	500	ETM	BAT
C_P	0.3190	0.2929	0.3024	**0.3360**	0.3258	0.2349	0.2509
NPMI	0.0641	0.0640	0.0672	**0.0748**	0.0642	0.0359	0.0487
UCI	0.4730	0.5150	**0.6998**	0.6624	0.5118	−0.0016	0.4192
Augmentation probability p							
Setting	5	10	15	20	25	ETM	BAT
C_P	**0.3181**	0.3172	0.3024	0.3161	0.3041	0.2349	0.2509
NPMI	0.0690	**0.0716**	0.0672	0.0679	0.0677	0.0359	0.0487
UCI	0.6574	0.6856	**0.6999**	0.5489	0.6003	−0.0016	0.4192

4.4. Human Evaluation

To assess the topic quality of the SSTM and baseline approaches, we also perform a human evaluation of the User Query dataset with 10 topic number settings. In detail, we recruit two experts to find the out-of-topic words, which are semantically inconsistent with

the topic, and count the number of irrelevant words. Finally, we average the counts and conduct the comparison in Table 7.

From the statistics, we could observe that topics extracted by the SSTM contain fewer out-of-topic words compared with other baselines. This further validates the superior performance of the SSTM in terms of topic coherence.

Table 7. Comparison results of human evaluation; a smaller number indicates better performance.

Model	LDA	WLDA	ATM	BAT	ETM	CNTM	BERTopic	CTMNeg	SSTM
Count	11.5	14	16.5	13.5	19.5	10.5	18	19.5	6

4.5. Case Study

We also carry out case studies and present concern examples in Figure 6 to validate the effectiveness of the SSTM. From the examples, we can observe that the SSTM uncovers concerns of Twitter users about "Student Cheating with ChatGPT" and "Impact on the Investment Market". Likewise, SSTM mines concerns about "Assitant on math" and "Solving Programming Puzzles using ChatGPT" from the Subreddit and User Query datasets.

Dataset	Concerns	Topic words and representative posts/queries
Twitter Posts	*Student cheating*	*Topic: school pass exam medical assignment assessment cheat classroom student teach* **username** @name How are high school teachers coping with students crushing their homework using #ChatGPT? Honestly, homework is dead.
	Surges of ChatGPT user	*Topic: reach million meeting user note month active premium monthly launch* **username** @name According to SimilarWeb, #ChatGPT has 75 million unique monthly users. If at least 10% of them sign up for the premium subscription:7.5 million x $42/month x 11 = $3.46 billion in revenue. At 60% profit margin, thats $2 billion in profit. In the first year.
	Impact on investment market	*Topic: product price investor market company billion financial investment public service* **username** @name #Microsoft Corp. is in talks to invest $10 billion into #OpenAI, the owner of #ChatGPT, which will value the #SanFrancisco-based firm at $29 billion, Semafor reported, citing people familiar with the matter.
Subreddit	*Assistant on math*	*Topic: skill problem learn tool math solve calculator calculation equation basic* **username** . 6 mo. ago This reminds of having a calculator but needing to do math in your head. I think it's important for kids to learn how to write and be creative.
	Humanization about ChatGPT	*Topic: robot experience conscious human thought consciousness matter physical emotion sentient* **username** . 7 mo. ago I Now if it was speaking to me in a human voice and was inside a little humanoid robot I might have some feelings tow
	Plugins of ChatGPT	*Topic: information engine search link internet browser chrome microsoft edge website* **username** . 6 mo. ago They are There is already a browser extension that puts a answer next to google search results
User Query	*ChatGPT for medical diagnosis*	*Topic: medical health effect exercise cause blood cell disease mental muscle* **Chatgpt user** In a clinical trial in spine surgery with fibrinogen, how would you show that blood loss, fibrinogen levels and FIBTEM clot amplitude are correlated since blood loss is used as a trigger,guided by FIBTEM clot amplitude?
	ChatGPT act as writing assistant and language tool	*Topic: language text sentence copy translate rewrite translation french spelling speak* **Chatgpt user** I want you to act as an French translator, spelling corrector and improver. I will speak to you in any language and you will detect the language, translate it and answer in the corrected and improved version of my text, in French. I want you to replace my simplified A0-level words and sentences with more beautiful and elegant, upper level French words and sentences. Keep the meaning same, but make them more literary. I want you to only reply the correction, the improvements and nothing else, do not write explanations
	ChatGPT for program-ming	*Topic: application code block framework database device implementation complete functionality request* **Chatgpt user** I got an out-of-memory error, can you tell me something about php memory optimization with code examples that actually Works?

Figure 6. Concern examples and representative posts/queries.

To prove the faithfulness of extracted concerns, we use Sentence-Bert [47] to design a retrieval scheme. Posts are selected according to the cosine similarity of contextualized embeddings of the topic words and text in the corpus.

We choose three distinct user concerns from three datasets and present the topic words along with corresponding posts or queries. The concerns in the Twitter Posts and Subreddit datasets primarily revolve around discussions regarding ChatGPT itself and its influence. On the other hand, the concerns in the User Query dataset highlight the specific tasks that users expect ChatGPT to address or complete. The above three results reveal the hot spots of public discussion on ChatGPT from various views.

4.6. Data Analysis

Due to the user privacy protection mechanism of social media platforms, we can not access user privacy information, such as age and gender. However, we conduct data analysis and present the differences in concerns extracted from different platforms (Twitter, Reddit, and Query). In this part, topic numbers are set to 10. And we present each mined concern with one keyphrase as shown in Table 8.

Table 8. Concerns mined by SSTM from Twitter Posts, Subreddit and User Query datasets.

Dataset	Summary
Twitter Posts	machine learning, student cheating, search engine, correct solution, music experience, marketing investment, medical information bias, dreams of freedom, surging ChatGPT user, digital marketing content
Subreddit	humanization about ChatGPT, assistant on math, search engine, tech industry cover, plugins of ChatGPT, content creation, engaging replies, academic writing tool, chatbot interaction, freedom and governance
User Query	medical diagnosis, market strategy, response protocol, academic writing tool, research and citation, programming assistant, customer engagement, character development, server application, blog post

It can be observed that several common concerns are discussed by users from different platforms, such as "*academic writing tool*" extracted from both Subreddit and User Query. Likewise, there are also some platform-specific user concerns like the "*programming assistant*" and "*student cheating*". These extracted concerns/topics present a comprehensive understanding of public opinion on ChatGPT.

4.7. Ablation Study

As the regularization terms and document representation approaches play key roles in our proposed SSTM, we conduct the corresponding ablation in this subsection to explore their impacts on topic modeling performance.

4.7.1. Regularization Term Ablation

To explore the roles of invariance, variance, and covariance regularization terms in the SSTM, we add one more dataset (Subreddit) in this subsection and conduct the ablation study on the Subreddit and User Query datasets with the topic number set to 10. The corresponding results are listed in Table 9; $SSTM_{-i}$, $SSTM_{-v}$ and $SSTM_{-c}$ mean variants that exclude invariance, variance and covariance terms.

Here, for a clearer comparison, we have retained the results of ETM and BAT, as their performance is quite competitive when the number of topics is set to 10. From the statistics, we observe that $SSTM_{-i}$ performs the worst among SSTM variants, and this may be attributed to the invariance objective having a high impact on topic quality. Similarly,

we observe that $SSTM_{-c}$ performs slightly worse than the SSTM. This discrepancy can be attributed to the covariance regularizer, which aids in decorrelating dimensions of the topic distribution and segregating irrelevant words into different topics for interpretable topic extraction. Additionally, we find a decline in the UTR metric for the $SSTM_{-v}$, which reveals that the variance regularizer has a positive impact on topic diversity. On the other hand, though these variants perform worse than the SSTM, they also surpass the competitive ETM on all the coherence and diversity metrics, which can be ascribed to the usage of the self-supervised learning framework and the Dirichlet prior.

Table 9. Comparison of ablation on regularization terms (invariance, variance, and covariance) utilized in SSTM.

	Subreddit				User Query			
Settings	**C_P**	**NPMI**	**UCI**	**UTR**	**C_P**	**NPMI**	**UCI**	**UTR**
$SSTM_{-i}$	0.2493	0.0531	0.4341	1.000	0.2947	0.0593	0.4495	1.000
$SSTM_{-v}$	0.2584	0.0555	0.4327	0.990	0.2965	0.0620	0.4353	0.990
$SSTM_{-c}$	0.2560	0.0563	0.4874	1.000	0.2970	0.0634	0.5021	1.000
$SSTM$	0.2719	0.0575	0.5309	1.000	0.3024	0.0672	0.6998	1.000
ETM	0.1388	0.0277	−0.1782	0.910	0.2349	0.0359	−0.0016	0.960
BAT	0.1901	0.0286	0.0655	0.780	0.2509	0.0487	0.4192	0.830

4.7.2. Document Representation Ablation

To examine if incorporating external semantic knowledge will improve modeling ability and explore the impact of the document representation approaches, we conduct this experiment using the pre-trained language model for representing documents on the Subreddit dataset. Concretely, for each document that contains the word sequence $\{w_1, w_2 \dots w_n\}$, we first generate the contextualized word representations $\{\vec{e_1}, \vec{e_2} \dots \vec{e_n}\}$. Then, the document is represented with the averaged word-level representations, computed with $\frac{1}{n}\sum_{i=1}^{n}\vec{e_i}$. The detailed comparison is listed in Table 10.

Table 10. Comparison of ablation on document representation approaches.

Settings	***K***	**C_P**	**NPMI**	**UCI**	**UTR**
$SSTM_P$	10	0.2982	0.0665	0.5488	1.000
	20	0.2550	0.0521	0.3661	1.000
	30	0.2418	0.0429	0.1936	0.970
	40	0.2421	0.0460	0.2335	0.950
$SSTM$	10	0.2719	0.0575	0.5309	1.000
	20	0.2374	0.0463	0.3368	1.000
	30	0.2261	0.0371	0.0410	0.967
	40	0.1933	0.0313	0.0519	0.922
LDA	10	0.2391	0.0460	−2.1031	0.920
	20	0.1931	0.0389	−2.1116	0.975
	30	0.1580	0.0283	−2.6291	0.960
	40	0.1364	0.0204	−2.5751	0.887

As shown in Table 10, the SSTM with the pre-trained language model $SSTM_P$ outperforms the SSTM and the competitive LDA across all the coherence and diversity metrics for all the topic number settings. For clarity, we only retain the LDA as the baseline as it performs the second best according to Table 3. Notably, the comparison proves that injecting external semantic knowledge can indeed improve the modeling ability of the SSTM.

5. Discussion

This study is based on the rapid development of the ChatGPT dialogue system, which has sparked extensive discussions and generated a large volume of text data. Such user-generated texts contain potential public expectations and concerns, and understanding these is crucial for the development of ChatGPT. We collect and process relevant data from three different social media platforms to mine public concerns. Although self-supervised learning has been widely applied in natural language processing and computer vision, there has been limited exploration of topic models using representation learning techniques. Thus, we formulate topic modeling as a representation learning task and propose the Self-Supervised Neural Topic Model (SSTM). To ensure topic quality and diversity, three training objectives (invariance, variance and covariance) are incorporated into the SSTM. The experimental results demonstrate the superior performance of the SSTM in discovering interpretable concerns/topics. Also, experimental results show that the SSTM could extract higher quality concerns/topics from the User Query dataset than those of Twitter Posts and Subreddit, which is caused by the data characteristics. Specifically, queries often contain more words than tweets and posts and have sufficient co-occurence information which is crucial for topic modeling and pattern discovery. Thus, we could draw the conclusion that our proposed SSTM is prone to mine higher-quality concerns/topics from platforms that can provide long-length texts.

Due to the effectiveness of the SSTM in mining public concerns, practical implications could be strengthened by the SSTM. For example, we could provide technical support for companies and policymakers to discover public concerns about other topics. Also, we will explore designing new policies (such as privacy enhancement measures and transparency reports) to address the ethical concerns about these AI technologies, especially in sensitive areas like education and healthcare.

6. Conclusions

In this paper, we propose the Self-Supervised Neural Topic Model, a novel neural topic modeling approach based on representation learning, for mining high-quality concerns about ChatGPT. In more detail, the SSTM utilizes the Dirichlet prior in the latent topic space to capture multiple semantic patterns. Additionally, to ensure the quality of the extracted concerns, it incorporates three regularization terms, which are invariance, variance, and covariance, for model training. The extensive experiments on three ChatGPT-related datasets show that our proposal could mine higher quality public concerns about this advanced dialogue system in terms of interpretability and diversity. As the proposed, the SSTM only employs the bag-of-words information, neglecting the sparse characteristics of the review text and the contextual information, and our future work will involve integrating external semantic knowledge (such as word embeddings and the pre-trained language models) into the modeling process and designing the knowledge-enhanced topic model for public concern extraction task.

Author Contributions: Conceptualization, R.W. and Z.H.; methodology, R.W. and X.L.; formal analysis and investigation, P.R. and S.C.; Writing—original draft preparation, X.L.; Writing—review and editing, P.R. and S.C.; funding acquisition, R.W. and Z.H.; visualization, X.L. and H.H.; supervision, G.S. and Z.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China (Nos. 62102192, 62162063), the Science and Technology Base and Talent Program of Guangxi under Grant number 2021AC19308 (Project contract number GuikeAD22035110), the Innovation and Entrepreneurship Program of Jiangsu Province (JSSCBS20210530), the fellowship of China Postdoctoral Science Foundation (2022M710071), and the Introduction of Talent Research and Research Fund of Nanjing

University of Posts and Telecommunications (NY220132), the Fundamental Research Funds for the Central Universities (aiia-24-01).

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: We extend our gratitude to the anonymous reviewers for their valuable feedback and beneficial recommendations. We strictly use the datasets referenced in the paper within the scope of their licenses, and we have removed personal information during the data preprocessing stage to protect privacy. Furthermore, we will only apply the data for academic research purposes.

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

Abbreviations

The following abbreviations are used in this manuscript:

SSTM	Self-Supervised Neural Topic Model
VAE	Variational Auto-Encoder
GANs	Generative Adversarial Networks
SSL	Self-Supervised Learning
NTMs	Neural Topic Models
LDA	Latent Dirichlet Allocation
NVDM	Neural Variational Document Mode
BAT	Bidirectional Adversarial-neural Topic

References

1. Liu, Y.; Han, T.; Ma, S.; Zhang, J.; Yang, Y.; Tian, J.; He, H.; Li, A.; He, M.; Liu, Z.; et al. Summary of ChatGPT/GPT-4 Research and Perspective Towards the Future of Large Language Models. *arXiv* **2023**, arXiv:2304.01852.
2. Abdullah, M.; Madain, A.; Jararweh, Y. ChatGPT: Fundamentals, Applications and Social Impacts. In Proceedings of the Ninth International Conference on Social Networks Analysis, Management and Security, SNAMS 2022, Milan, Italy, 29 November–1 December 2022; IEEE: Piscataway, HJ, USA, 2022; pp. 1–8. [CrossRef]
3. Zheng, S.; Yahya, Z.; Wang, L.; Zhang, R.; Hoshyar, A.N. Multiheaded deep learning chatbot for increasing production and marketing. *Inf. Process. Manag.* **2023**, *60*, 103446. [CrossRef]
4. Hill-Yardin, E.; Hutchinson, M.; Laycock, R.; Spencer, S. A Chat (GPT) about the future of scientific publishing. *Brain Behav. Immun.* **2023**, *110*, 152–154. [CrossRef] [PubMed]
5. Li, Z.; Hu, H.; Wang, H.; Cai, L.; Zhang, H.; Zhang, K. Why does the president tweet this? Discovering reasons and contexts for politicians' tweets from news articles. *Inf. Process. Manag.* **2022**, *59*, 102892. [CrossRef]
6. Zhao, H.; Phung, D.Q.; Huynh, V.; Jin, Y.; Du, L.; Buntine, W.L. Topic Modelling Meets Deep Neural Networks: A Survey. In Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Montreal, QC, Canada, 19–27 August 2021; pp. 4713–4720. [CrossRef]
7. Nguyen, T.; Luu, A.T. Contrastive Learning for Neural Topic Model. In *Advances in Neural Information Processing Systems 34, Proceedings of the Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, Virtual, 6–14 December 2021*; Neural Information Processing Systems Foundation, Inc.: La Jolla, CA, USA, 2021; pp. 11974–11986.
8. Wang, R.; Zhou, D.; He, Y. Open Event Extraction from Online Text using a Generative Adversarial Network. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 November 2019; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 282–291. [CrossRef]
9. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. In Proceedings of the 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, 14–16 April 2014.
10. Wang, R.; Hu, X.; Zhou, D.; He, Y.; Xiong, Y.; Ye, C.; Xu, H. Neural Topic Modeling with Bidirectional Adversarial Training. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, 5–10 July 2020; Association for Computational Linguistics: Stroudsburg, PA, USA, 2020; pp. 340–350. [CrossRef]
11. Adhya, S.; Lahiri, A.; Sanyal, D.K.; Das, P.P. Improving Contextualized Topic Models with Negative Sampling. In Proceedings of the ICON, International Conference on Cognitive Neuroscience, Barcelona, Spain, 24–28 April 2022; Association for Computational Linguistics: Stroudsburg, PA, USA, 2022; pp. 128–138.

12. Dieng, A.B.; Ruiz, F.J.R.; Blei, D.M. Topic Modeling in Embedding Spaces. *Trans. Assoc. Comput. Linguist.* **2020**, *8*, 439–453. [CrossRef]
13. Grootendorst, M. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv* **2022**, arXiv:2203.05794.
14. Wang, R.; Zhou, D.; He, Y. ATM: Adversarial-neural Topic Model. *Inf. Process. Manag.* **2019**, *56*, 102098. [CrossRef]
15. Zhou, X.; Bu, J.; Zhou, S.; Yu, Z.; Zhao, J.; Yan, X. Improving topic disentanglement via contrastive learning. *Inf. Process. Manag.* **2023**, *60*, 103164. [CrossRef]
16. Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; Deny, S. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In Proceedings of the 38th International Conference on Machine Learning, ICML 2021, Virtual, 18–24 July 2021; Volume 139, pp. 12310–12320.
17. Misra, I.; van der Maaten, L. Self-Supervised Learning of Pretext-Invariant Representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 6706–6716. [CrossRef]
18. Hendrycks, D.; Mazeika, M.; Kadavath, S.; Song, D. Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty. In Proceedings of the Neural Information Processing Systems NeurIPS, Vancouver, BC, Canada, 8–14 December 2019; pp. 15637–15648.
19. Bijoy, M.B.; Pebbeti, B.P.; Akondi, S.M.; Fathaah, S.A.; Raut, A.; Pournami, P.N.; Jayaraj, P.B. Deep Cleaner—A Few Shot Image Dataset Cleaner Using Supervised Contrastive Learning. *IEEE Access* **2023**, *11*, 18727–18738. [CrossRef]
20. Fini, E.; Astolfi, P.; Alahari, K.; Alameda-Pineda, X.; Mairal, J.; Nabi, M.; Ricci, E. Semi-supervised learning made simple with self-supervised clustering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, 17–24 June 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 3187–3197. [CrossRef]
21. Bardes, A.; Ponce, J.; LeCun, Y. VICRegL: Self-Supervised Learning of Local Visual Features. In *Advances in Neural Information Processing Systems 35, Proceedings of the Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, 28 November–9 December 2022*; Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A., Eds.; Neural Information Processing Systems Foundation, Inc.: La Jolla, CA, USA, 2022.
22. Yan, Y.; Li, R.; Wang, S.; Zhang, F.; Wu, W.; Xu, W. ConSERT: A Contrastive Framework for Self-Supervised Sentence Representation Transfer. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Bangkok, Thailand, 1–6 August 2021; Volume 1, pp. 5065–5075.
23. Zhu, H.; Zheng, Z.; Soleymani, M.; Nevatia, R. Self-Supervised Learning for Sentiment Analysis via Image-Text Matching. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Singapore, 23–27 May 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 1710–1714. [CrossRef]
24. Mu, L.; Jin, P.; Zhang, Y.; Zhong, H.; Zhao, J. Synonym recognition from short texts: A self-supervised learning approach. *Expert Syst. Appl.* **2023**, *224*, 119966. [CrossRef]
25. Luo, X.; Ju, W.; Gu, Y.; Mao, Z.; Liu, L.; Yuan, Y.; Zhang, M. Self-supervised Graph-level Representation Learning with Adversarial Contrastive Learning. *ACM Trans. Knowl. Discov. Data* **2024**, *18*, 1–23. [CrossRef]
26. Yi, S.; Ju, W.; Qin, Y.; Luo, X.; Liu, L.; Zhou, Y.; Zhang, M. Redundancy-Free Self-Supervised Relational Learning for Graph Clustering. *arXiv* **2023**, arXiv:2309.04694. [CrossRef]
27. Ju, W.; Wang, Y.; Qin, Y.; Mao, Z.; Xiao, Z.; Luo, J.; Yang, J.; Gu, Y.; Wang, D.; Long, Q.; et al. Towards Graph Contrastive Learning: A Survey and Beyond. *arXiv* **2024**, arXiv:2405.11868. [CrossRef]
28. Ju, W.; Gu, Y.; Mao, Z.; Qiao, Z.; Qin, Y.; Luo, X.; Xiong, H.; Zhang, M. GPS: Graph Contrastive Learning via Multi-scale Augmented Views from Adversarial Pooling. *arXiv* **2024**, arXiv:2401.16011. [CrossRef]
29. Yang, J.; Xu, H.; Mirzoyan, S.; Chen, T.; Liu, Z.; Liu, Z.; Ju, W.; Liu, L.; Xiao, Z.; Zhang, M.; et al. Poisoning medical knowledge using large language models. *Nat. Mac. Intell.* **2024**, *6*, 1156–1168. [CrossRef]
30. Huang, J.; Chen, L.; Guo, T.; Zeng, F.; Zhao, Y.; Wu, B.; Yuan, Y.; Zhao, H.; Guo, Z.; Zhang, Y.; et al. MMEvalPro: Calibrating Multimodal Benchmarks Towards Trustworthy and Efficient Evaluation. *arXiv* **2023**, arXiv:2407.00468.
31. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* **2003**, *3*, 993–1022.
32. Miao, Y.; Yu, L.; Blunsom, P. Neural Variational Inference for Text Processing. In Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, 19–24 June 2016; Volume 48, pp. 1727–1736.
33. Srivastava, A.; Sutton, C. Autoencoding Variational Inference For Topic Models. In Proceedings of the 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017.
34. Miyato, T.; Kataoka, T.; Koyama, M.; Yoshida, Y. Spectral Normalization for Generative Adversarial Networks. In Proceedings of the 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, 30 April–3 May 2018.
35. Howard, A.; Pang, R.; Adam, H.; Le, Q.V.; Sandler, M.; Chen, B.; Wang, W.; Chen, L.; Tan, M.; Chu, G.; et al. Searching for MobileNetV3. In Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Republic of Korea, 27 October–2 November 2019; IEEE: Piscataway, HJ, USA, 2019; pp. 1314–1324. [CrossRef]

36. Wallach, H.M.; Mimno, D.M.; McCallum, A. Rethinking LDA: Why Priors Matter. In *Advances in Neural Information Processing Systems 22, Proceedings of the 23rd Annual Conference on Neural Information Processing Systems 2009, Vancouver, BC, Canada, 7–10 December 2009*; Curran Associates, Inc.: Newry, UK, 2009; pp. 1973–1981.
37. Gretton, A.; Borgwardt, K.M.; Rasch, M.J.; Schölkopf, B.; Smola, A.J. *A Kernel Method for the Two-Sample-Problem*; MIT Press: Cambridge, MA, USA, 2006; pp. 513–520.
38. Baktashmotlagh, M.; Harandi, M.T.; Salzmann, M. Distribution-Matching Embedding for Visual Domain Adaptation. *J. Mach. Learn. Res.* **2016**, *17*, 1–30.
39. Berlinet, A.; Thomas-Agnan, C. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2011.
40. Lafferty, J.D.; Lebanon, G. Diffusion Kernels on Statistical Manifolds. *J. Mach. Learn. Res.* **2005**, *6*, 129–163.
41. Bardes, A.; Ponce, J.; Lecun, Y. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. In Proceedings of the ICLR 2022—International Conference on Learning Representations, Virtual, 25–29 April 2022.
42. Miao, Y.; Grefenstette, E.; Blunsom, P. Discovering Discrete Latent Topics with Neural Variational Inference. In Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017; PMLR: Breckenridge, CO, USA, 2017; Volume 70, pp. 2410–2419.
43. Tolstikhin, I.O.; Bousquet, O.; Gelly, S.; Schölkopf, B. Wasserstein Auto-Encoders. In Proceedings of the 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, 30 April–3 May 2018.
44. Röder, M.; Both, A.; Hinneburg, A. Exploring the Space of Topic Coherence Measures. In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM 2015, Shanghai, China, 2–6 February 2015; ACM: New York, NY, USA, 2015; pp. 399–408. [CrossRef]
45. Chang, J.D.; Boyd-Graber, J.L.; Gerrish, S.; Wang, C.; Blei, D.M. Reading Tea Leaves: How Humans Interpret Topic Models. In *Advances in Neural Information Processing Systems 22, Proceedings of the 23rd Annual Conference on Neural Information Processing Systems 2009, Vancouver, BC, Canada, 7–10 December 2009*; Curran Associates, Inc.: Newry, UK, 2009; pp. 288–296.
46. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015.
47. Reimers, N.; Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 November 2019; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 3980–3990. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.