

Generative AI has lowered the barriers to computational social sciences

Yongjun Zhang¹

¹ Department of Sociology and Institute for Advanced Computational Science, Stony Brook University, Stony Brook, New York, United States

©Version 1, Updated on Nov 9, 2023

* yongjun.zhang@stonybrook.edu

Abstract

Generative artificial intelligence (AI) has revolutionized the field of computational social science, unleashing new possibilities for analyzing multimodal data, especially for scholars who may not have extensive programming expertise. This breakthrough carries profound implications for the realm of social sciences. Firstly, generative AI can significantly enhance the productivity of social scientists by automating the generation, annotation, and debugging of code. Secondly, it empowers researchers to delve into sophisticated data analysis through the innovative use of prompt engineering. Lastly, the educational sphere of computational social science stands to benefit immensely from these tools, given their exceptional ability to annotate and elucidate complex codes for learners, thereby simplifying the learning process and making the technology more accessible.

Introduction

The swift advancement of generative AI, particularly with OpenAI’s groundbreaking release of GPT-4 turbocharged models, has significantly streamlined the computational social science (CSS) landscape. Traditionally, CSS scholars grappled with the intricate task of distilling valuable insights from diverse and multimodal data sources — text, imagery, audio, and video. The pre-generative AI era demanded from them a mastery of programming to adeptly fine-tune pre-trained large language or computer vision models using transfer learning for sophisticated downstream tasks like sentiment analysis and image classification.

Now, the vanguard of OpenAI’s suite — GPT-4V, Whisper V3, and advanced Text-to-Speech (TTS) technologies — empowers scholars to seamlessly convert varied inputs into coherent text outputs. Even more transformative is the capacity for nuanced prompt engineering, enabling a code interpreter that simplifies the generation, annotation, and debugging of code, thereby making complex computational tasks more accessible.

Within the scope of this commentary, I strive to illuminate how generative AI tools are dismantling the traditional barriers faced by digital humanities and CSS practitioners, particularly those with limited coding proficiency. Following this introduction, Section 2

will elucidate the process by which scholars can harness generative AI for coding purposes — generating, annotating, and debugging with unprecedented ease. Subsequently, Section 3 will delve into the direct analysis of multimodal data using GPT-4 models, exploring the burgeoning implications this holds for AI’s trajectory within social science research.

Coding is no longer a big challenge to computational social scientists

A significant hurdle for social scientists delving into computational research has been the necessity to craft sophisticated code capable of managing complex operations, such as distilling pivotal information from textual and visual data. In the era preceding the advent of generative AI, academics would typically fine-tune pre-trained large language and vision models — like BERT, ResNet, and Swin Transformer — to perform nuanced text and image classification tasks. For example, Y. Zhang, Lin, Wang, and Fan (2023) adeptly honed the Robustly Optimized Bidirectional Encoder Representations from Transformers (RoBERTa) model to analyze the sentiments and topics of 25 million tweets, uncovering prevalent sinophobia within the Chinese language discourse amid the initial stages of the pandemic. Similarly, H. Zhang and Pan (2019) meticulously calibrated a substantial convolutional neural network, previously trained on ImageNet, combined with a recurrent neural network to categorize images and texts from Weibo, thereby compiling a database chronicling collective action in China. Such deep learning endeavors traditionally require fluency in frameworks like PyTorch or TensorFlow, posing a formidable barrier to novices at the confluence of social movements and computational social sciences.

Generating Codes

A pivotal advancement in the realm of generative AI has been its capacity to craft functional code for end-users. Tools such as GitHub Copilot, Code LLaMa, and the ChatGPT code interpreter have become staples for scholars and engineers, facilitating daily code generation. This proves especially advantageous for social scientists whose expertise in programming may be nascent. I will demonstrate this through OpenAI’s ChatGPT (GPT-4), highlighting its code interpreter’s ability to generate scripts that fine-tune a Swin Transformer model for image classification in the context of protests.

To illustrate, consider my project where I amassed a dataset of approximately 1.6 million images from around the globe, potentially depicting social protests. These images reside on my network-attached storage, organized within a directory named *images_by_country*, further segmented into subdirectories labeled by country codes. Additionally, I have leveraged the UCLA protest image dataset, which comprises 40,764 images — 11,659 of which are protest-related — each tagged with various visual attributes and sentiments (Won, Steinert-Threlkeld, & Joo, 2017). It’s noteworthy that while the images in the training dataset are sorted into ‘train’ and ‘test’ folders, the corresponding labels are housed in distinct ‘train-annotation’ and ‘test-annotation’ TSV files. These files contain ‘fname’ and ‘protest’ labels, where ‘fname’ matches the image filenames in the train and test directories. The core objectives are two-fold: firstly, to script Python code that loads the UCLA protest image dataset

to refine a pre-trained Swin Transformer v2 model for protest image identification, and secondly, to apply the refined model to classify the 1.6 million image dataset. The ensuing sections will narrate the iterative prompting process undertaken to fulfill these tasks.

Prompt: Write pytorch code to fine tune a swin transformer model v2.
Training images are stored in train and test folders while labels are stored in two separate tsv files train_annot.txt and test_annot.txt. Note that these txt files contain variables like fname and protest. fname is file names corresponding to images in train and test folders. You need to write python code to load image and label data and then finetune, evaluate, and save swin transformer model. Finally, you need to use the saved model to infer images stored in images_by_country with subfolders named by countries. Make sure you write python codes to check whether these iamges files are valid.

ChatGPT: Below is a Python code to fine-tune a Swin Transformer model using PyTorch. This code assumes you have the necessary environment and the Swin Transformer pre-trained model installed.

The steps are as follows:

Read the annotation files and prepare the datasets.
Define a custom dataset class to load images and labels.
Define the model, loss function, and optimizer.
Train the model.
Evaluate the model on the test dataset.
Save the fine-tuned model.
Infer on new images from images_by_country and check for valid image files.

First, make sure you have installed the required libraries:

```
pip install torch torchvision timm pandas pillow
```

Here's the complete Python code:

```
import os
import pandas as pd
from PIL import Image
from torchvision import transforms
from torch.utils.data import Dataset, DataLoader
from torch import nn, optim
import torch
import timm
from torchvision.io import read_image
from torchvision.transforms.functional import to_pil_image
from torchvision.utils import save_image

# 1. Prepare the datasets
train_annot = pd.read_csv('train_annot.txt', sep='\t')
test_annot = pd.read_csv('test_annot.txt', sep='\t')

# 2. Define a custom dataset class
class ProtestDataset(Dataset):
    def __init__(self, annotations_file, img_dir, transform=None):
        self.img_labels = pd.read_csv(annotations_file, sep='\t')
        self.img_dir = img_dir
        self.transform = transform

    def __len__(self):
        return len(self.img_labels)

    def __getitem__(self, idx):
        img_path = os.path.join(self.img_dir, self.img_labels.iloc[idx, 0])
        image = Image.open(img_path).convert('RGB')
        label = self.img_labels.iloc[idx, 1]
```

```

        if self.transform:
            image = self.transform(image)
        return image, label

# Define the transform
transform = transforms.Compose([
    transforms.Resize((224, 224)),
    transforms.ToTensor(),
    transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]),
])

# Create the datasets
train_dataset = ProtestDataset('train_annot.txt', 'train', transform=transform)
test_dataset = ProtestDataset('test_annot.txt', 'test', transform=transform)

# 3. Define the model, loss function, and optimizer
device = torch.device('cuda' if torch.cuda.is_available() else 'cpu')

model = timm.create_model('swin_v2', pretrained=True)
model.head = nn.Linear(model.head.in_features, 1) # Adjust for binary classification

model = model.to(device)
criterion = nn.BCEWithLogitsLoss()
optimizer = optim.Adam(model.parameters(), lr=0.001)

# 4. Train the model
def train(dataloader, model, loss_fn, optimizer):
    size = len(dataloader.dataset)
    for batch, (X, y) in enumerate(dataloader):
        X, y = X.to(device), y.to(device)

        # Compute prediction error
        pred = model(X)
        loss = loss_fn(pred, y.unsqueeze(1).float())

        # Backpropagation
        optimizer.zero_grad()
        loss.backward()
        optimizer.step()

    if batch % 100 == 0:
        loss, current = loss.item(), batch * len(X)
        print(f"loss: {loss:>7f} [{current:>5d}/{size:>5d}]")

# 5. Evaluate the model
def test(dataloader, model, loss_fn):
    size = len(dataloader.dataset)
    model.eval()
    test_loss, correct = 0, 0
    with torch.no_grad():
        for X, y in dataloader:
            X, y = X.to(device), y.to(device)
            pred = model(X)
            test_loss += loss_fn(pred, y.unsqueeze(1).float()).item()
            correct += (pred.sigmoid().round() == y).type(torch.float).sum().item()
    test_loss /= size
    correct /= size
    print(f"Test Error: \n Accuracy: {(100*correct):>0.1f}%, Avg loss: {test_loss:>8f} \n")

# Dataloaders
train_dataloader = DataLoader(train_dataset, batch_size=64, shuffle=True)
test_dataloader = DataLoader(test_dataset, batch_size=64, shuffle=False)

# Training loop
epochs = 10
for t in range(epochs):
    print(f"Epoch {t+1}\n-----")
    train(train_dataloader, model, criterion, optimizer)

```

```

    test(test_dataloader, model, criterion)
print("Done!")

# 6. Save the model
torch.save(model.state_dict(), 'swin_transformer_finetuned.pth')

# 7. Infer on new images and check for valid image files
def is_valid_image(file_path):
    try:
        with Image.open(file_path) as img:
            img.verify() # verify that it is, in fact, an image
        return True
    except (IOError, SyntaxError) as e:
        print('Bad file:', file_path)
        return False

model.eval()
countries_dir = 'images_by_country'
countries = [d for d in os.listdir(countries_dir) if os.path.isdir(os.path.join(countries_dir, d))]

for country in countries:
    country_dir = os.path.join(countries_dir, country)
    for image_file in os.listdir(country_dir):
        image_path = os.path.join(country_dir, image_file)
        if not is_valid_image(image_path):
            continue
        image = Image.open(image_path).convert('RGB')
        image = transform(image).unsqueeze(0).to(device)
        with torch.no_grad():
            prediction = model(image).sigmoid().item()
        print(f"Image {image_file} in {country} predicted as
        {'Protest' if prediction > 0.5 else 'Non-Protest'} with confidence {prediction}")

```

Annotating Codes

The initial foray into code generation with AI tools may yield code snippets that are complex or not immediately operational, especially for those just embarking on their coding journey. Understanding the intricacies of the functions and the specifics of parameter setting can be daunting. To bridge this gap, one can enlist the help of a code interpreter to meticulously dissect and elucidate the code line by line. This guided exploration facilitates a deeper comprehension of the technical nuances. Here's an example of how one might engage ChatGPT to perform such detailed code annotations:

"I've received this block of code from the initial output, but as someone new to programming, I'm struggling to grasp its structure and logic. Could you walk me through this code, providing annotations for each segment and explaining the purpose of the different parameters? This way, I can better understand what each part of the code is intended to do and how I might adjust it to fit my specific needs."

By posing the request in this manner, you are likely to receive a breakdown that not only sheds light on the code's current functionality but also imparts knowledge that can empower you to modify and adapt the code independently in the future.

Prompt: I don't understand the train function in the above generated code. Please annotate the code line by line and explain it to me in detail. Make sure you explain the parameters in those functions.

Certainly! Here is a detailed annotation of the train function, breaking down each line and explaining what it does:

```

def train(dataloader, model, loss_fn, optimizer):
    """
    The train function is responsible for one epoch of training the model.

    Parameters:
    - dataloader: An instance of DataLoader which provides batches of the dataset.
    - model: The neural network model that is being trained.
    - loss_fn: The loss function that is used to evaluate the performance of the model.
    - optimizer: The optimization algorithm used to update the weights of the model.
    """

    # Size of the dataset being used for training.
    size = len(dataloader.dataset)
    # Set the model into training mode which enables the training specific operations like dropout.
    model.train()

    # Loop through each batch in the DataLoader.
    for batch, (X, y) in enumerate(dataloader):
        # X holds the batch of images, y holds the corresponding labels.

        # Move the input and target batches to the same device as the model (GPU or CPU).
        X, y = X.to(device), y.to(device)

        # Perform a forward pass: compute the predicted outputs by passing inputs to the model.
        pred = model(X)

        # Compute the loss by comparing the predictions to the actual labels.
        # The unsqueeze operation is used to add a dimension to y so it matches the shape of pred.
        loss = loss_fn(pred, y.unsqueeze(1).type_as(pred))

        # Before the backward pass, use the optimizer to zero out any gradients from the previous step
        # to prevent accumulation from previous forward passes.
        optimizer.zero_grad()

        # Perform a backward pass: compute the gradient of the loss with respect to model parameters.
        loss.backward()

        # After calculating the gradients, we update the parameters with the optimizer step.
        optimizer.step()

        # Print out the loss every 100 batches to check the training progress.
        # This is not strictly necessary but helpful to know that something is happening.
        if batch % 100 == 0:
            loss, current = loss.item(), batch * len(X)
            print(f"loss: {loss:>7f} [{current:>5d}/{size:>5d}]")

```

The function train takes four parameters:

dataloader: This parameter should be a DataLoader object that provides batches of data in the form of (input_data, labels). It abstracts the complexity of iterating over the dataset, and automatically handles the generation of minibatches and shuffling of the data if specified.

model: The neural network model that you are training. This should be an instance of a PyTorch neural network model which inherits from nn.Module.

loss_fn: A loss function that takes the model's predictions and the true labels and calculates a numerical value representing how well the model is performing. In PyTorch, this is typically an instance of a class that inherits from nn.modules.loss._Loss.

optimizer: This is an instance of an optimizer class from torch.optim. It is responsible for updating the model parameters with the gradients computed during backpropagation.

In the function body, model.train() is called to put the model in training mode.

This is important because some types of layers, like dropout and batch normalization, behave differently during training than during inference. 310
311
312

During the loop, the model makes predictions on the input batch X and then the loss is calculated by comparing these predictions to the true labels y . After that, the gradients are computed by `loss.backward()`, and the optimizer updates the model's parameters with `optimizer.step()`. Finally, some logging is performed to track the progress. 313
314
315
316
317

Note: The `.type_as(pred)` function is used to make sure that y is of the same data type as `pred`, which is necessary for the loss function to work properly. 318
319
This might be needed because sometimes labels could be in a different type (like `int`) which is not compatible with the model predictions (like `float`). 320
321
322

Remember, `device` should be defined outside this function and should represent the computing device (CPU or GPU) where the tensors are to be placed. 323
324

Debugging Codes 325

When the initial code output from ChatGPT doesn't align with your requirements, it's a prudent step to iterate with more specific prompts. This iterative process is essential for refining the code to match your precise needs. Here's how you might proceed: 326
327
328

1. Identify Specific Issues: Start by pinpointing the specific areas where the code falls short. Is it a matter of functionality, efficiency, or compatibility with your existing systems? 329
330
331
2. Craft Targeted Prompts: Once you've identified the issues, formulate targeted prompts that directly address these points. For example, if the code is not functioning as expected, you might ask, "How can I modify this function to achieve [desired outcome]?" 332
333
334
3. Provide Context: Give ChatGPT as much context as possible. This could include the desired end result, constraints you're working within, or even the steps you want the code to follow. 335
336
337
4. Ask for Explanations: If you don't understand the code, ask for comments or documentation that explain what each part does. For instance: "Can you add comments to the code explaining each function and parameter?" 338
339
340
5. Iterate as Needed: Sometimes, it might take several rounds of back-and-forth to get the code right. Don't hesitate to keep asking for adjustments until the code meets your specifications. 341
342
343
6. Test the Code: After receiving the updated code, test it in your environment. If there are still issues, provide feedback on the results and ask for further revisions. 344
345

Here's an example of a refined prompt for further code development: 346

Prompt 1: Revise the above code using transformer library from huggingface and make sure you use `trainer` function instead of `pytorch` code. 347
348

Prompt 2: Revise the above code you generated and make it parallel when inferring new image dataset. Also update the image transform using `swin transformer preprocessor`. 349
350
351

```

from transformers import SwinForImageClassification, SwinProcessor, Trainer, TrainingArguments 352
from datasets import load_metric 353
from torch.utils.data import DataLoader 354
from PIL import Image 355
import pandas as pd 356
import torch 357
import os 358
from concurrent.futures import ThreadPoolExecutor 359

# Custom dataset class 360
class ProtestDataset(torch.utils.data.Dataset): 361
    def __init__(self, img_dir, annotations_file, processor=None): 362
        self.img_labels = pd.read_csv(annotations_file, delimiter='\t') 363
        self.img_dir = img_dir 364
        self.processor = processor 365

    def __len__(self): 366
        return len(self.img_labels) 367

    def __getitem__(self, idx): 368
        img_path = os.path.join(self.img_dir, self.img_labels.iloc[idx, 0]) 369
        image = Image.open(img_path).convert("RGB") 370
        label = self.img_labels.iloc[idx, 1] 371
        if self.processor: 372
            processed = self.processor(images=image, return_tensors="pt") 373
            image = processed.pixel_values.squeeze() # Apply preprocessing 374
            return {'pixel_values': image, 'labels': torch.tensor(label)} 375

# Load the processor and model 376
processor = SwinProcessor.from_pretrained('microsoft/swin-base-patch4-window7-224') 377
model = SwinForImageClassification.from_pretrained('microsoft/swin-base-patch4-window7-224', num_labels=2) 378

# Prepare datasets 379
train_dataset = ProtestDataset(img_dir='train', annotations_file='train_annot.txt', processor=processor) 380
eval_dataset = ProtestDataset(img_dir='test', annotations_file='test_annot.txt', processor=processor) 381

# Define the training arguments 382
training_args = TrainingArguments( 383
    output_dir='./results', 384
    evaluation_strategy="epoch", 385
    learning_rate=2e-5, 386
    per_device_train_batch_size=8, 387
    per_device_eval_batch_size=16, 388
    num_train_epochs=3, 389
    weight_decay=0.01, 390
) 391

# Initialize the Trainer 392
trainer = Trainer( 393
    model=model, 394
    args=training_args, 395
    train_dataset=train_dataset, 396
    eval_dataset=eval_dataset, 397
    compute_metrics=load_metric("accuracy"), 398
) 399

# Train and evaluate 400
trainer.train() 401
trainer.evaluate() 402

# Save the fine-tuned model 403
trainer.save_model("swin_transformer_finetuned") 404

# Helper function to check image validity 405
def is_valid_image(file_path): 406
    try: 407
        Image.open(file_path).verify() 408
        return True 409
    except Exception: 410
        return False 411

# Helper function to process and predict an image 412
def process_and_predict(image_path, processor, model): 413
    if not is_valid_image(image_path): 414
        return image_path, None 415
    image = Image.open(image_path).convert('RGB') 416
    inputs = processor(images=image, return_tensors="pt").pixel_values 417
    inputs = inputs.to(model.device) 418
    with torch.no_grad(): 419
        predictions = model(inputs).logits.argmax(-1).cpu().numpy() 420
    return image_path, predictions[0] 421

# Load the trained model for inference 422
model = SwinForImageClassification.from_pretrained("swin_transformer_finetuned").to("cuda") 423

# Perform parallel inference on images from different countries 424
def run_inference_in_parallel(directory, processor, model): 425
    image_paths = [os.path.join(directory, fname) for fname in os.listdir(directory) if is_valid_image(os.path.join(directory, fname))] 426
    predictions = {} 427

    with ThreadPoolExecutor(max_workers=os.cpu_count()) as executor: 428
        futures = {executor.submit(process_and_predict, image_path, processor, model): image_path for image_path in image_paths} 429

```



```

    for future in futures:
        image_path, prediction = future.result()
        predictions[image_path] = prediction

    return predictions

# Assuming images_by_country is a directory with subdirectories named by countries
countries_dir = 'images_by_country'
all_predictions = {}

for country in os.listdir(countries_dir):
    country_dir = os.path.join(countries_dir, country)
    if not os.path.isdir(country_dir):
        continue
    predictions = run_inference_in_parallel(country_dir, processor, model)
    all_predictions[country] = predictions

# Print or otherwise utilize all_predictions as needed
for country, predictions in all_predictions.items():
    print(f"Predictions for {country}: {predictions}")

```

I ran the above codes using a workstation with two RTX 3090 GPUs and 128 core CPUs. It returns the following error.

```

Code bug:
    ImportError: cannot import name 'SwinProcessor' from 'transformers'

```

I use the following prompt to ask ChatGPT to fix the code for me.

```

Prompt: After ran your code, I got the following error:
    ImportError: cannot import name 'SwinProcessor' from 'transformers'.
    Fix it.

ChatGPT: I apologize for the confusion earlier.
    As of my last update, the transformers library does not include a SwinProcessor.
    Instead, you would use a feature extractor designed for image models.
    Let's correct the code by using SwinForImageClassification with
    an appropriate feature extractor.
...
# Load the feature extractor
feature_extractor = AutoFeatureExtractor.from_pretrained('microsoft/swin-base-patch4-window7-224')
...

```

After a meticulous process of iterative debugging, the refined Swin Transformer model impressively attained an accuracy rate of 94%. Similarly, I built a ResNet50 model which also demonstrated robust performance, achieving an 85% accuracy rate. The polished codes that underpinned the training and inference of these images are comprehensively detailed in the Appendix.

Generative AI can help social scientists process multi-modal data instead of fine-tuning large language and vision models.

The preceding illustration underscores the potential of generative AI tools in streamlining the coding process, particularly in fine-tuning sophisticated pre-trained vision models for image classification. Although I maintain that coding has ceased to be a formidable obstacle for engaging in computational research, it is still imperative for social scientists to possess a foundational grasp of these innovative tools.

Beyond code facilitation, generative AI also boasts the capability to directly analyze multimodal data, encompassing text, imagery, and audio. To exemplify this, I turn to the enhanced GPT-4 model with vision. Consider the image of the Women’s March in Washington, sourced from NPR, which captures marchers united in their call for women’s rights. By uploading this image to GPT-4 and crafting targeted prompts, I was able to efficiently extract pertinent information. This instance not only demonstrates the prowess of generative AI in processing diverse datasets but also its utility as a powerful tool for social scientists to distill insights without deep technical expertise.

Prompt 1: Describe the image using concise sentences.
Then tell me if it is protest-related and show me your reasoning.

ChatGPT: The image shows a group of people actively participating in a march. They are holding a large banner that reads "END WOMEN'S WAGE DISC." and "2019 WOMEN'S MARCH ON WASHINGTON", indicating the event is the 2019 Women's March in Washington D.C. The participants are carrying signs with various slogans and a few are wearing pink hats, which are symbols associated with the Women's March movement. One person in the foreground is using a megaphone, suggesting speeches or chants are part of this event.

This is definitely protest-related, as the Women's March is a well-known event advocating for women's rights and gender equality. The presence of signs, a large banner with a clear message regarding women's wage disparity, and the collective action of marching are consistent with protest activities.

Prompt 1: Summarize the image and tell me the sentiment, content, and crowd size.

ChatGPT: The image captures a moment from the 2019 Women's March in Washington D.C., with participants holding a banner advocating against women's wage discrimination. The sentiment appears to be one of solidarity and activism, aimed at promoting gender equality and social justice. The signs and symbols, including pink hats and raised fists, communicate a demand for change and unity among the marchers.

Content-wise, the photo includes a diversity of participants of various ages and ethnicities, and some of the visible signs carry messages of empowerment and resistance.

The crowd size in the immediate vicinity of the banner is substantial, with dozens of people visible. However, the photograph does not show the entire scope of the march, so the total number of participants is likely much larger, potentially in the thousands, which is typical for such events.

For scholars aiming to process a multitude of images, the integration of OpenAI’s GPT-4 with its vision API emerges as a formidable tool. While access to this service is contingent upon a subscription, the investment unlocks GPT-4’s sophisticated data analytics capabilities. Social scientists can harness these features for elementary data analysis tasks, thereby enhancing their research productivity significantly.

Take, for instance, my utilization of GPT-4 to dissect the Dynamics of Collective Action dataset, culminating in a detailed map that delineates the geographical disparities in protest occurrences throughout the United States. For an in-depth understanding of the DoCA dataset, visit this link. Upon uploading a segment of this dataset to GPT-4 and prompting it appropriately, I could extract intricate data visualizations. Notably, within the realm of data governance, users have the autonomy to restrict GPT-4 from employing their data in model training, ensuring control and confidentiality.

Prompt: Describe the CSV file and summarize the data pattern.

ChatGPT:
I'll first need to load and examine the contents of the CSV file to provide you with a description and summarize the data pattern. Let me do that for you.

```
import pandas as pd

# Load the CSV file into a DataFrame
file_path = '/mnt/data/data.csv'
data_df = pd.read_csv(file_path)

# Display the first few rows of the DataFrame to understand its structure
data_df.head()
```

The CSV file contains a dataset with the following columns:

rptyy: Appears to represent a year, with floating-point numbers (likely intended to be integers).

eventid: An identifier for an event, also numeric.

state1: A text field containing abbreviations of US states.

city1: A text field containing city names.

The data pattern suggests that this file is a record of events, with each row representing a specific event identified by eventid, occurring in a particular city1 and state1 during a given rptyy (year).

To summarize the data pattern further, I can provide statistics like the number of unique years, event counts, and the frequency of events by state and city. Would you like me to do that?

Prompt: Yes. Please provide the frequency of events by state.

ChatGPT:

The frequency of events by state is as follows:

New York (NY): 8,613 events
 Washington D.C. (DC): 1,902 events
 California (CA): 1,398 events
 New Jersey (NJ): 1,371 events
 Massachusetts (MA): 881 events
 Illinois (IL): 772 events

Prompt: Write R code using GGPILOT2 to visualize the geographic distribution of protest events across states in the US.

ChatGPT:

Certainly! Below is an example of R code using ggplot2 to visualize the geographic distribution of protest events across states in the U.S. The code assumes you have a data frame named data with a column state1 for state abbreviations and that each row represents an event. This code will not run in this environment, but you can run it in your local R environment.

```

library(ggplot2)
library(maps)
library(dplyr)

# Ensure the data is in the correct format (factor) for the state abbreviations
data$state1 <- as.factor(data$state1)

# Count the number of events by state
state_event_counts <- data %>%
  group_by(state1) %>%
  summarise(count = n())

# Merge the counts with the map data
us_states_map <- map_data("state")
state_event_map <- merge(us_states_map, state_event_counts, by.x = "region", by.y = "state1", all.x = TRUE)
state_event_map <- state_event_map[order(state_event_map$order), ]

# Replace NA with 0 for states with no events
state_event_map$count[is.na(state_event_map$count)] <- 0

# Plot the map
ggplot(state_event_map, aes(x = long, y = lat, group = group, fill = count)) +
  geom_polygon(color = "white") +
  expand_limits(x = us_states_map$long, y = us_states_map$lat) +
  scale_fill_continuous(low = "white", high = "red", name = "Event Count") +
  labs(title = "Geographic Distribution of Protest Events Across US States") +
  theme_minimal() +
  theme(legend.position = "right",
        panel.border = element_blank(),
        panel.grid.major = element_blank(),
        panel.grid.minor = element_blank(),
        axis.text = element_blank(),
        axis.title = element_blank(),
        axis.ticks = element_blank())

```

This code performs the following tasks:

1. Loads the required libraries (ggplot2, maps, and dplyr).
2. Converts the state abbreviations to a factor, which is necessary for mapping.
3. Counts the number of events per state using dplyr's group_by and summarise functions.
4. Merges the event counts with the map data of US states.
5. Handles states with no events by setting their counts to 0.
6. Uses ggplot2 to create a choropleth map, filling the states based on the number of events. The geom_polygon function is used to create the map, and the fill aesthetic is based on the event count.
7. Adjusts the color scale from white (for no events) to red (for many events).
8. Adds a title and adjusts the theme for better visualization.

The Implications to AI for Social Sciences

651

In this commentary, I have endeavored to illustrate how the advent of generative AI has democratized computational methodologies, enabling social scientists, even those with minimal coding proficiency, to harness the power of cutting-edge computational tools. The crux of this technological leap is that coding, once a formidable barrier to entry in the digital humanities and computational social science, is being redefined as a facilitator rather than a gatekeeper.

652
653
654
655
656
657

Generative AI is a game-changer for the social sciences, as posited by (Bail, n.d.) and (Davidson, 2023). Its primary benefit lies in alleviating the coding burden, thereby unlocking higher productivity and efficiency among researchers. This liberation from the intricacies of code generation allows social scientists to invest more deeply in the conceptualization of their research, potentially fostering the advancement of social theories. Moreover, the transformative influence of generative AI has reshaped the research paradigm, making computational tools accessible to a broader spectrum of scholars, thereby catalyzing discoveries in social science. Significantly, these tools have also proven to be exceptional assets in pedagogy, offering substantial capabilities to annotate and clarify coding concepts for students in computational social science disciplines.

658
659
660
661
662
663
664
665
666
667

However, this progress is not without its caveats and conditions. Despite the argument that coding should not pose a substantial hurdle, scholars must still possess fundamental skills to assess both code and the quality of its outputs. For example, when generating code for models like the Swin Transformer, a foundational understanding of deep learning principles is essential. Thankfully, a basic course in computational social science is typically sufficient to surmount this learning curve. A more pressing concern is the evaluation of outputs, particularly when employing models such as GPT-4 for vision and advanced data analysis tasks. For those dealing with extensive image datasets, establishing protocols to systematically verify the accuracy of the results is imperative and an area ripe for scholarly inquiry.

668
669
670
671
672
673
674
675
676
677

While some of the tools referenced here are proprietary, budgetary constraints need not stifle scholarly ambition. Researchers with a modicum of coding knowledge can turn to open-source alternatives like LLaMa2, which do require some computing resources to execute code locally. For those without access to significant computational power, free resources such as Google Colab present a viable solution, thereby continuing to lower the barriers and broaden the horizons for research in the social sciences.

678
679
680
681
682
683

Acknowledgments

684

I acknowledge the support from the Institute for Advanced Computational Science for access to high performance computing systems and OpenAI APIs and ChatGPT. This version of draft v2 is also revised and refined by ChatGPT.

685
686
687
688

Appendix

688

With several rounds of iterations, I managed to get the following useful codes to fine-tune
swin transformer model and use it to infer new dataset. These scripts were split into two.

689

690

```
from urllib.request import urlopen
import ssl

ssl._create_default_https_context = ssl._create_unverified_context

import os
import pandas as pd
import torch
from torch.utils.data import Dataset, DataLoader
from torchvision.io import read_image
from transformers import SwinV2ForImageClassification, SwinV2Config, AutoImageProcessor, Trainer, TrainingArguments
from transformers import DataCollatorWithPadding
from torch.utils.data.dataloader import default_collate
import numpy as np # Make sure numpy is imported

# Custom dataset class remains unchanged
class CustomImageDataset(Dataset):
    def __init__(self, annotations_file, img_dir, transform=None, target_transform=None):
        self.img_labels = pd.read_csv(annotations_file, sep="\t")[['fname', 'protest']]
        self.img_dir = img_dir
        self.transform = transform
        self.target_transform = target_transform

    def __len__(self):
        return len(self.img_labels)

    def __getitem__(self, idx):
        img_path = os.path.join(self.img_dir, self.img_labels.iloc[idx, 0])
        image = read_image(img_path).float() # Convert to float for ensuring float32 precision

        # Ensure image is in CHW format expected by PyTorch models
        if image.ndim != 3:
            raise ValueError(f"Image at {img_path} does not have three dimensions. Found {image.ndim} dimensions.")

        label = int(self.img_labels.iloc[idx, 1])

        # Apply the image processor. It handles necessary image transformations
        processed_images = self.transform(image, return_tensors="pt")

        # SwinV2 model expects keys 'pixel_values' for images. Also, we don't need to unsqueeze here as 'return_tensors' takes care of it
        return {
            'pixel_values': processed_images['pixel_values'].squeeze(0), # remove the added batch dimension
            'labels': torch.tensor(label) # Ensure label is a tensor
        }

def custom_data_collator(features):
    # 'features' is a list of dicts with the outputs of your dataset's __getitem__ method
    # You need to aggregate the different elements of these dicts here, for instance by creating a new dict with aggregated tensors

    batch = {}

    # Aggregate the 'pixel_values' and 'labels' from each feature in the batch
    batch['pixel_values'] = torch.stack([f['pixel_values'] for f in features])
    batch['labels'] = torch.tensor([f['labels'] for f in features])

    return batch

# Load the Swin Transformer V2 model from HuggingFace
model_name = "microsoft/swinv2-base-patch4-window8-256"
config = SwinV2Config.from_pretrained(model_name)
config.num_labels = 2 # Set the number of output classes

model = SwinV2ForImageClassification.from_pretrained(model_name, config=config, ignore_mismatched_sizes=True)

# Initialize the AutoImageProcessor
processor = AutoImageProcessor.from_pretrained(model_name)

# Create dataset instances
train_dataset = CustomImageDataset(annotations_file='./annot_train.txt', img_dir='./img/train', transform=processor)
test_dataset = CustomImageDataset(annotations_file='./annot_test.txt', img_dir='./img/test', transform=processor)

# Define the training arguments
training_args = TrainingArguments(
    output_dir='./results', # Output directory for model predictions and checkpoints
    num_train_epochs=3, # Number of training epochs
    per_device_train_batch_size=128, # Batch size for training
    per_device_eval_batch_size=128, # Batch size for evaluation
```

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

```

warmup_steps=500,          # Number of warmup steps for learning rate scheduler      774
weight_decay=0.01,        # Strength of weight decay                                                       775
logging_dir='./logs',     # Directory for storing logs                                                      776
logging_steps=10,         # Log saving interval                                                             777
evaluation_strategy="epoch", # Evaluate at the end of each epoch                                              778
save_strategy="epoch",    # Save the model at the end of each epoch                                        779
load_best_model_at_end=True, # Load the best model when finished training (based on evaluation)              780
)                                                                    781

# Define a simple function to compute metrics (e.g., accuracy)                782
def compute_metrics(eval_pred):                                           783
    logits, labels = eval_pred                                           784
    # Convert logits to a tensor if it's a numpy array                      785
    if isinstance(logits, np.ndarray):                                     786
        logits = torch.tensor(logits)                                     787
    # Similarly, ensure labels are a tensor; this is not in your error but might cause issues later 788
    if isinstance(labels, np.ndarray):                                    789
        labels = torch.tensor(labels)                                     790
    predictions = torch.argmax(logits, dim=-1) # Using 'dim' as the argument key for clarity 791
    accuracy = (predictions == labels).float().mean()                    792
    return {'accuracy': accuracy.item()}                                   793
                                                                            794
# Initialize the Trainer with the custom data collator                      795
trainer = Trainer(                                                       796
    model=model,                                                         797
    args=training_args,                                                 798
    train_dataset=train_dataset,                                         799
    eval_dataset=test_dataset,                                           800
    compute_metrics=compute_metrics,                                     801
    data_collator=custom_data_collator, # Use your custom data collator 802
)                                                                           803
                                                                            804
# Start training                                                           805
trainer.train()                                                         806
                                                                            807
# Evaluate the model                                                       808
results = trainer.evaluate()                                             809
                                                                            810
# Save the model                                                           811
trainer.save_model("./swin_v2_b_model_protect")                          812
                                                                            813
# You can also get the model's predictions on the test dataset using the following code: 814
# predictions = trainer.predict(test_dataset)                             815
                                                                            816
import os                                                                817
import pandas as pd                                                     818
import torch                                                            819
from torchvision.transforms import functional as F                       820
from transformers import SwinV2ForImageClassification, AutoImageProcessor 821
from concurrent.futures import ThreadPoolExecutor                       822
from PIL import Image                                                  823
from tqdm import tqdm                                                  824

# Load the fine-tuned model and processor                                825
model_dir = "./swin_v2_b_model_protect"                                  826
model_name = "microsoft/swinv2-base-patch4-window8-256"                827
model = SwinV2ForImageClassification.from_pretrained(model_dir)         828
processor = AutoImageProcessor.from_pretrained(model_name)               829
                                                                            830
# Define a function to predict the class of a single image                831
def predict_image(image_info):                                           832
    country_folder, image_file = image_info                               833
    image_path = os.path.join(image_folder, country_folder, image_file) 834
    try:                                                                  835
        # Attempt to open the file as an image                             836
        image = Image.open(image_path)                                    837
        # Ensure image is in RGB format                                     838
        if image.mode != 'RGB':                                           839
            image = image.convert('RGB')                                    840
        # Convert PIL image to PyTorch tensor                              841
        image = F.to_tensor(image).float()                                  842
        inputs = processor(image, return_tensors="pt", do_rescale=False) 843
        outputs = model(**inputs)                                          844
        prediction = torch.argmax(outputs.logits, dim=-1)                 845
        return {"Country": country_folder, "Image": image_file, "Prediction": prediction.item()} 846
    except Exception as e:                                                847
        print(f"Error processing file {image_path}: {str(e)}")           848
    return None                                                           849

# Iterate through the files in the "images_by_country" folder and make predictions 850
image_folder = "images_by_country"                                       851
image_infos = []                                                         852

```

```

for country_folder in os.listdir(image_folder):
    country_path = os.path.join(image_folder, country_folder)
    if os.path.isdir(country_path):
        for image_file in os.listdir(country_path):
            image_infos.append((country_folder, image_file))

# Use ThreadPoolExecutor to parallelize the image prediction
with ThreadPoolExecutor() as executor:
    predictions = list(tqdm(executor.map(predict_image, image_infos), total=len(image_infos), desc="Predicting"))

# Remove None entries (if any)
predictions = [pred for pred in predictions if pred is not None]

# Save the predictions to a CSV file
predictions_df = pd.DataFrame(predictions)
predictions_df.to_csv("predictions_by_country.csv", index=False)

```

References

- Bail, C. A. (n.d.). Can generative ai improve social science?
- Davidson, T. (2023). Start generating: Harnessing generative artificial intelligence for sociological research.
- Won, D., Steinert-Threlkeld, Z. C., & Joo, J. (2017). Protest activity detection and perceived violence estimation from social media images. In *Proceedings of the 25th acm international conference on multimedia* (pp. 786–794).
- Zhang, H., & Pan, J. (2019). Casm: A deep-learning approach for identifying collective action events with text and image data from social media. *Sociological Methodology*, *49*(1), 1–57.
- Zhang, Y., Lin, H., Wang, Y., & Fan, X. (2023). Sinophobia was popular in chinese language communities on twitter during the early covid-19 pandemic. *Humanities and Social Sciences Communications*, *10*(1), 1–12.