

## Research Article

# NewsHomepages: Homepage Layouts Capture Information Prioritization Decisions

Ben Welsh<sup>1</sup>, Naitian Zhou<sup>2</sup>, Alexander Spangher<sup>3</sup>

1. Thomson Reuters (United Kingdom), London, United Kingdom; 2. University of California, Berkeley, United States; 3. University of Southern California, United States

Information prioritization plays an important role in how humans perceive and understand the world. Homepage layouts serve as a tangible proxy for this prioritization. In this work, we present NewsHomepages, a large dataset of over 3,000 new website homepages (including local, national and topic-specific outlets) captured twice daily over a three-year period. We develop models to perform pairwise comparisons between news items to infer their relative significance. To illustrate that modeling organizational hierarchies has broader implications, we applied our models to rank-order a collection of local city council policies passed over a ten-year period in San Francisco, assessing their “newsworthiness”. Our findings lay the groundwork for leveraging implicit organizational cues to deepen our understanding of information prioritization.

Ben Welsh, Arda Kaz, and Alexander Spangher equally contributed to this work.

Corresponding author: Alexander Spangher, [spangher@usc.edu](mailto:spangher@usc.edu)

## 1. Introduction

The way humans prioritize and organize information is central to attention and understanding. Prior work has explored the need for effective prioritization to manage information overload<sup>[1]</sup> and including case-studies of prioritization approaches<sup>[2][3]</sup>. However, larger scale analyses, specifically with an aim towards predictive modeling, are often limited by lack of available data and methods.

In this work, we seek to rectify this by focusing on a prominent source of publicly available, highly curated examples of prioritization decisions. News organizations’ homepages are meticulously crafted

by professional human editors: their layouts offer manifestations of the news organization's prioritization principles<sup>[4]</sup>. We use homepage layouts to study information organization at scale. We present NewsHomepages, to our knowledge the first large-scale homepage layout dataset, consisting of over 363k full page screenshots, HTML, extract links and other metadata from over 3,000 websites, which we captured twice daily over a three-year period, from local, magazine, and national news homepages. With this dataset, we ask: *how well can we model the editorial decisions behind homepage layout? What do these decisions tell us about the priorities of the newspaper and the world?*

To answer these questions, we develop models to process and learn from our dataset at scale. First, we develop a novel bootstrapping approach to robustly parse homepage screen captures and determine precise spatial positions of news articles on homepages. Then, we train models to perform pairwise comparisons between articles based on their positions within layouts. By interpreting positional cues as indicators of significance, our models infer the relative importance of information.

The ability to learn these decisions with high accuracy indicates the robust signal this dataset captures. We conducted two downstream experiments utilizing these models; in each, we applied our pairwise comparison models to rank-order a list of text. In the first experiment, we use models trained on the newsworthy judgments of one news outlet to sort the news articles of another. We find surprising nuances: for instance, despite *Breitbart* (a right-leaning outlet) being topically dissimilar to *Mother Jones* (a left-leaning outlet), their newsworthy preferences are among the most correlated of outlets we studied. In the second analysis, we used each outlet's newsworthiness judgments to rank-order local city council policies passed in San Francisco<sup>[5]</sup>. By assessing these policies' "newsworthiness" through the lens of organizational prominence, we highlight how implicit cues in information structure can transfer between domains and serve as useful tools (in this case, to help human journalists find relevant story-topics). Our contributions are as follows:

- We introduce NewsHomepages, large-scale dataset of homepage layouts from over 3,000 local, national and topic-specific news organizations, and develop a robust, weakly-supervised method to parse articles.
- We demonstrate that editorial principles behind information prioritization can be learned, by training pairwise comparison models to predict article size and positionality.

We show via two case-studies – (1) newsworthiness agreement between outlets, and (2) generating newsworthiness rankings for non-news corpora – that such learned principles can generalize beyond

the corpora we study and provide useful tools for end-users.

This work opens new avenues for exploring how implicit cues in digital environments reflect latent priorities and organizational principles. By modeling these cues, we can gain richer insights into the mechanisms that influence human perception.

## 2. Background and Dataset

Visual cues for editorial importance on homepages have a deep history in the design principles of physical newspapers<sup>[6]</sup> and result from deliberate editorial processes. At *The New York Times*, for example, top editors and designers convened daily in the renowned *Page One* meeting<sup>[7]</sup> to determine the most important articles for the print newspaper the next day<sup>1</sup>. In the digital era, meetings like this evolved into *Homepage Meetings*<sup>[8]</sup>, influencing the design and content placement on the website's homepage for the upcoming day. As such, homepages continue to be distillations of professional judgement and priorities.

One visual cue editors use is **positional placement**, with articles positioned towards the top and left of a page considered more important<sup>[9]</sup>. This stems from observations that readers naturally begin scanning from the top-left corner<sup>[10]</sup>. Secondly, the **space** articles occupy is considered: larger articles or headlines are perceived as more important<sup>[11]</sup>. In print media, prominence is conveyed through more column space; in digital media, longer headlines, featured images, and extended summaries. Finally, **graphics and design** also play a pivotal role in signaling the importance of news stories. Articles accompanied by photographs, videos, or other multimedia elements are often deemed more significant<sup>[12]</sup>. The use of capital letters, bold fonts, and color further enhances a story's prominence.

We find few large-scale computational analyses studying these attributes. To enable a more precise study of editorial judgement, we construct a large corpus of news homepage layouts, over which we can track these indicators of relative importance.

### 2.1. Compilation of News Homepages

We compiled a list of 3,489 news homepages, as of the time of this writing, which we scraped twice daily on an ongoing basis over a period of three years. From 2019–2024, we have collected a total of 363,340 total snapshots. Our dataset collection is actively maintained and facilitated by a large contributing community of over 35 activists, developers and journalists. We collect homepages from

national news outlets (e.g., *The New York Times*, *The Wall Street Journal*), state-level news outlets (e.g., *San Francisco Chronicle*, *Miami Herald*), as well as local and subject-matter-specific news sources. Table 1 provides a sample of the different categories of news homepages included in our dataset, and a full list can be found in the appendix. Additionally, we collect homepages from news websites of over 32 countries in 17 languages (please see Tables 6 and 8 for a more detailed breakdown). This is an ongoing and expanding effort: we encourage contributors to add their own news homepages of interest using for our suite of tools to scrape.<sup>2</sup> We hope to further diversify the news sources in the dataset that we collect.

Category	Example Outlets
National	The New York Times, The Wall Street Journal, NPR, Bloomberg
State-level	San Francisco Chronicle, Miami Herald, Chicago Tribune
Local	Sturgis-Journal, The Daily Jeffersonian, LAist, The Desert Sun
Subject-specific	The Weather Channel, Chessbase, ESPN
International	India Today, Ukrinform, BBC, Prensa Grafica, Japan Times

**Table 1.** Sample of News Homepages by Category

## 2.2. Data Collection Pipeline

Our dataset collection runs in a cron job twice a day, and uploads data to Internet Archive. For each snapshot, we store the following information:

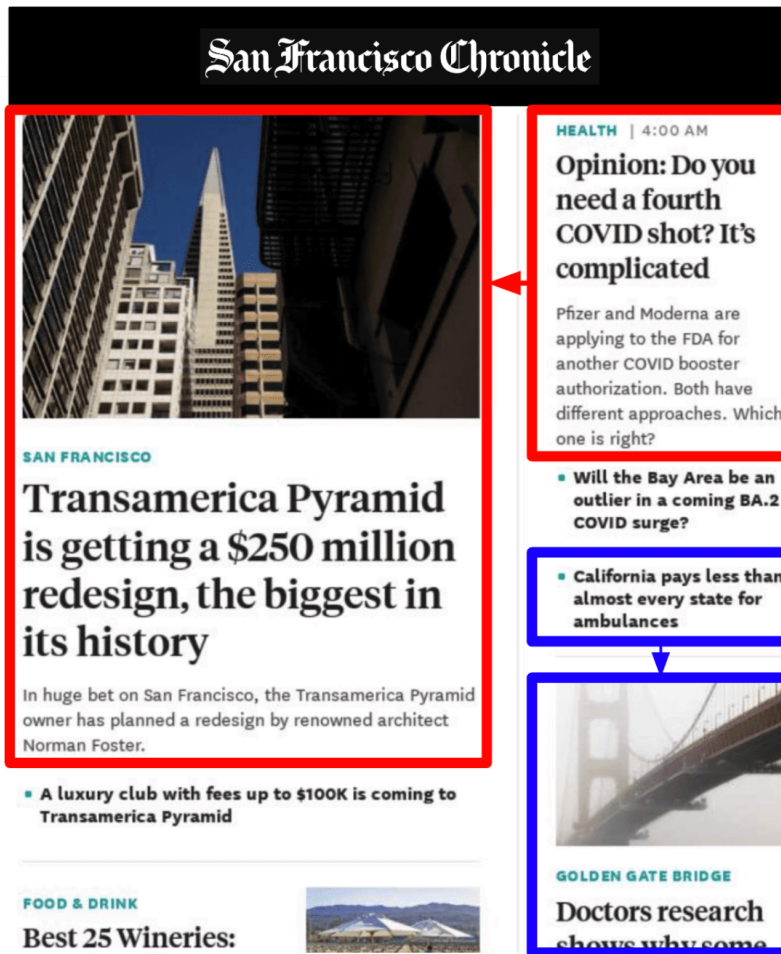
1. **All links on the page:** We store a flat-list of hyperlinks on every homepage and associated text.
2. **Full-page screenshots:** We store JPGs of each complete homepage as we render it.

**Complete HTML snapshots (subset of pages):** For a subset of homepages, we save a compressed version of the webpage, including all CSS files and images, using SingleFile.<sup>3</sup>

In addition to our Internet Archive storage,<sup>4</sup> we also synchronize with Wayback Machine to store these homepages, providing a secondary backup and ensuring long-term preservation.

### 3. Dataset Processing

In order to robustly extract visual attributes for each article on a homepage (i.e. size, position, presence of graphics), we need to determine bounding boxes for all articles on a homepage. Examples of bounding boxes are shown in Figure 1: each bounding box, also referred to as *article card*, covers all information directly associated with that article.



**Figure 1.** Two “newsworthiness” signals that editors make to guide reader attention are shown above. (1) **Position** (i.e. articles that are placed above, ↑, and left, ← relative to other articles are more important \newcitehays2018analysis). (2) **Size** (i.e. articles that are larger than other articles are more important) (3) **Graphics and Font** (i.e. articles with graphics and images are more important). We release NewsHomepages, a large dataset of over 3,000 homepages, collected twice-daily over three years, study information prioritization in this setting. We show can model these decisions at scale and demonstrate the usefulness of these models on two downstream tasks.

Layout parsing is a well-researched field<sup>[13][14]</sup>. However, homepages present unique challenges due to their diverse structures: text of varying size, fonts, colors and images are easily perceived by

humans. However, because none of the largest supervised datasets<sup>[15]</sup> are specific to our task, we find that existing resources fail for parsing homepages. So, we bootstrap a supervised detection task.

### 3.1. Bootstrapping a Bounding Box Detector

Following other bootstrapping approaches<sup>[16]</sup>, we: (1) develop a simple deterministic algorithm to generate candidate data, (2) apply a filtering step to exclude low-quality data, (3) use our high-precision dataset to train a more robust classifier. Figure 7, in the Appendix, provides an overview of the pipeline.

#### *Step 1: Find Bounding Boxes Deterministically*

We design a deterministic algorithm, called the DOM-Tree algorithm, to start our bootstrapping process. At a high level, the algorithm traces each `<a>` tag in the Document Object Model (DOM) and extracts the largest subtree in the DOM that contains *only a single `<a>` tag* (illustrated in Figure 4, Appendix). This method can extract the maximal bounding box for each article, however it faces robustness challenges, for example, if a link exists *within* an article card (e.g. a link to an authors page, as shown in Figure 4b, Appendix.)

We apply this algorithm to a subset of the NewsHomepages dataset, combining 15 homepages each from all outlets for which we have HTML files, JPEG snapshots, and hyperlink json files (approximately 15,000 homepages). Since each outlet typically maintains a consistent layout on their homepages across samples, we include more outlets for generalizability.

		FP #1	FP #2	FN #1	FN #2	Total Errors	% Correct
<b>Challenge dataset</b>	DOM-Tree algorithm	117	137	127	265	646	61.3%
	Detectron2 Model	25	23	27	87	162	90.3%
<b>Clean dataset</b>	DOM-Tree algorithm	12	20	0	13	45	97.1%
	Detectron2 Model	15	24	0	18	57	96.3%

**Table 2.** Error analysis of bounding box detection methods comparing the DOM-Tree algorithm and a Detectron2 model across two datasets: the challenge dataset and the clean dataset. The challenge dataset is formed by selecting the bottom 10% of articles based on the match between OCR-extracted text and retrieved link text (described in Section 3 Step 2), while the clean dataset contains well-matched articles. Error types are divided into false positives (FP #1: multiple articles in one box, FP #2: no articles in a box) and false negatives (FN #1: partially captured articles, FN #2: articles not captured). As can be seen, the our trained model performs at par on the DOM-Tree algorithm in the clean settings and is far more robust in noisy settings.

### *Step 2: Filter Low-Quality Bounding Box Extractions*

We take several filtering steps to prevent “drift”<sup>[16]</sup>. (1) First, we exclude non-news article links (e.g. log-in pages) by training simple text classifier to distinguish between URLs to news articles and others. We manually labeled over 2,000 URLs. The model achieves an accuracy of 96%. (2) Then, we exclude bounding boxes that did not overlap highly with link text. We determine this by first rendering the HTML pages as images and overlaying bounding boxes, then running OCR to extract the bounding-box text. (3) Finally, we exclude bounding boxes with improperly rendered images<sup>5</sup>. To address this, we again rendered HTML pages as an image and employed the YOLO object detection model<sup>[17]</sup> to compare these images to the JPEGs in our archive. If a screenshot was not within 80% of the detection count of the archived snapshot, we discarded the snapshot. Overall, this multi-stage filtering process significantly reduced the number of boxes that did not correspond to actual articles and removed many websites that contained broken or corrupt data, enhancing the quality of the training data.



### *Step 3: Train a More Robust Classifier*

Now with our dataset in hand, we trained a Detectron2 model<sup>[18]</sup>. Our model uses ResNet-101 as a backbone with a Feature Pyramid Network (FPN) for extracting multi-scale features and Smooth L1 loss for bounding box regression. During training, we used a base learning rate of 0.02 with a linear warmup over the first 1000 steps. We trained the model for 10,000 steps in total, with learning rate reductions after 5000 steps. A weight decay of 0.0001 and momentum of 0.9 were also employed. The training ran on  $4 \times A40$  GPUs for 24 hours.

### *3.2. Evaluation and Results*

To evaluate the quality of our bounding box detection, we conducted manual validation for four types of errors: 1) bounding boxes that contain multiple articles, 2) bounding boxes that contain no articles, 3) bounding boxes missing parts of an article, and 4) articles that are not captured.

We used the OCR text-matching method, as described in Section 3, to identify particularly challenging homepages.<sup>6</sup> We compared errors on the cleanest 10% of homepages (Clean) and the least-clean 10% (Challenge). As shown in Table 2, our computer vision model (Detectron2) significantly improved the accuracy of bounding box detection in contexts where the DOM-Tree algorithm struggles (the Detectron2 model had an Card Correct % score of 90.3% while the DOM-Tree had a score of 61.3%). For Clean pages, the Detectron2 model performed similarly to the DOM-Tree algorithm, with error differences being minimal and both models achieving high accuracy (above 96%). The combination of deterministic algorithms and machine learning techniques allow us to achieve a more robust extraction of article attributes from diverse homepage layouts.

	Size			Position x Size		
Model Name	F1	Precision	Recall	F1	Precision	Recall
Flan-t5-base	91.9	91.8	92.0	70.7	80.5	63.1
Flan-t5-Large	66.6	51.1	95.4	54.9	43.4	74.7
Roberta Base	91.0	91.7	90.4	64.9	78.1	55.6
Roberta Large	85.4	85.9	84.9	47.2	74.4	34.6
Distilbert-Base-Uncased	93.1	93.3	92.9	75.2	81.8	69.6

**Table 3.** Performance metrics on NYTimes data for different models

Outlet	Accuracy	F1	Recall	Prec.
phoenixluc	57.1	70.3	57.4	90.7
newsobserver	75.0	72.5	74.3	70.7
slate	72.4	61.6	66.2	57.7
jaxdotcom	75.2	63.4	65.5	61.4
arstechnica	64.7	17.5	41.4	11.1
airwaysmagazine	72.5	73.7	78.9	69.1
denverpost	73.7	67.8	70.5	65.4
thedailyclimate	82.0	80.9	81.3	80.6
breitbartnews	68.9	22.8	54.7	14.4
foxnews	67.3	38.6	55.6	29.5
motherjones	71.4	63.0	68.7	58.2
thehill	68.8	55.5	59.8	51.7
wsj	70.0	48.0	52.0	44.6

**Table 4.** Performance metrics on a sampling of outlets, including on ones we used for the downstream experiments in Section 5.

## 4. Homepage Modeling

Given precise layout information for the 363k homepages in our dataset, we arrive again at the core question of this research: how regular and predictable is an outlet’s layout decisions? We initially hypothesized that modeling homepage placement would be challenging. As shown in Figures 3, 5, 6<sup>7</sup> and described by<sup>[19]</sup>, certain areas of many homepages lack clear editorial consistency. This introduces noise and makes it difficult to learn a uniform policy guiding article placement decisions. Further, learning a single set of policies is challenged by the changing news cycle; some days have lots of news while others have less.

#### 4.1. Modeling Approach

A homepage is intended to present a collection of articles as a cohesive bundle; individual articles do not exist in isolation<sup>[20]</sup>. Predicting the placement of a single article without considering the context of other articles would be overly noisy and potentially ineffective<sup>[21]</sup>. Conversely, attempting to predict the placement of all articles simultaneously poses a combinatorial challenge that is computationally infeasible.

To address this issue, we formulate our modeling task as a pairwise preference problem. Specifically, we consider pairs of articles  $(a_1, a_2)$  and train models to predict a binary preference variable  $p$ , where  $p_o(a_1 > a_2) = 1$  if article  $a_1$  is preferred over article  $a_2$  for outlet  $o$ , and  $p_o(a_1 > a_2) = 0$  otherwise.

We explore three variations of preference criteria for the preference variable,  $p$ :

1. **Size-based Preference:** We define  $p_o(a_1 > a_2) = 1$  if article  $a_1$  occupies more surface area on the homepage than article  $a_2$ , assuming that prominent articles are given more space<sup>[22]</sup>.
2. **Position-based Preference:** We set  $p_o(a_1 > a_2) = 1$  if article  $a_1$  is placed in a more favorable location on the homepage than article  $a_2$ , such as higher up or more to the left, based on common reading patterns<sup>[23]</sup>.
3. **Combined Size and Position Preference:** Here,  $p_o(a_1 > a_2) = 1$  if article  $a_1$  either occupies more surface area or is in a more favorable position than article  $a_2$ , particularly focusing on articles that are in the top 10.

To model these preference variables,  $p$ , we train a simple Transformer-based binary classifier, `distilbert-base(X)`, which classifies a text sequence  $X$ . Our model concatenates the input articles:  $X = a_1 < \text{sep} > a_2$  as input; the model learns to recognize the `<sep>` token as a boundary between the first and the second articles.

#### 4.2. Modeling Variations

We explored different modeling variations on the *New York Times* homepages, as they have a variety of content, shares and functionalities on their site<sup>[24]</sup>. We test 5 different models: {`distilbert-base-uncased`, `flan-t5-base`, `flan-t5-large`, `roberta-base`, `roberta-large`} and constructed a training dataset of 74,857 article-pairs and a test dataset consisting of 18,715 datapoints consisting of pairs of NYTimes articles from same homepages.

We experienced exploding gradients in the `flan-t5-large` and `RoBERTa-large` models, motivating us to use a learning rate limit of  $5e-5$  for all the models, for the sake of equal comparison. We applied Parameter-Efficient-Fine-Tuning<sup>[25]</sup> on `flan-t5-base`, `flan-t5-large`, `roberta-base`, `roberta-large` models to minimize overfitting, as we had limited of datapoints.

The `distilbert-base-uncased` model outperforms other models (Table 3). We trained on 4xA40 GPUs and 16xA100 GPUs, and implemented gradient clipping after observing gradient explosion.

### 4.3. Dataset Selection and Processing

From our list of 3,000 outlets, we select 31 outlets for detailed analysis. We selected well-known outlets in various categories, including different political leanings (left-leaning vs. right-leaning<sup>8</sup>), local and national levels, and varied subject matters such as science, chess and aviation. For each outlet, we collected between 200 and 300 homepage snapshots, resulting in 1,000 to 50,000 pairs of articles. We created an 80/20 train/test split and trained `distilbert-base-uncased` models for each outlet. We trained each model with  $5e-5$  learning rate limit, 3 epochs, 0.01 weight decay.

Each article in our dataset includes its textual representation as it appeared on the homepage. To enhance the reliability of our models, we undertake several data processing steps informed by preliminary experiments: (1) we only sample pairs of articles that are adjacent on the homepage, to curate preference pairs that are more likely to be challenging and topically similar. Secondly, we clean the textual representations by stripping out any times, dates, and formatting elements. We also remove author names to prevent the models from learning biases based on authors who might be favored by the organization. Please refer to Appendix A for a detailed list of the outlets used and the specific number of data points associated with each.

### 4.4. Results

We show our results in Table 4. While some models (e.g. `Breitbart`) perform noticeably poorly, we note that the majority of our models score above  $f_1 > .6$ . We do not find a significant correlation between model performance and training set size. We were surprised to observe the tractability of this task; this indicates that many of the concerns we had about noise were either handled by our preprocessing steps, or not as important as we believed.

## 5. Demonstrations

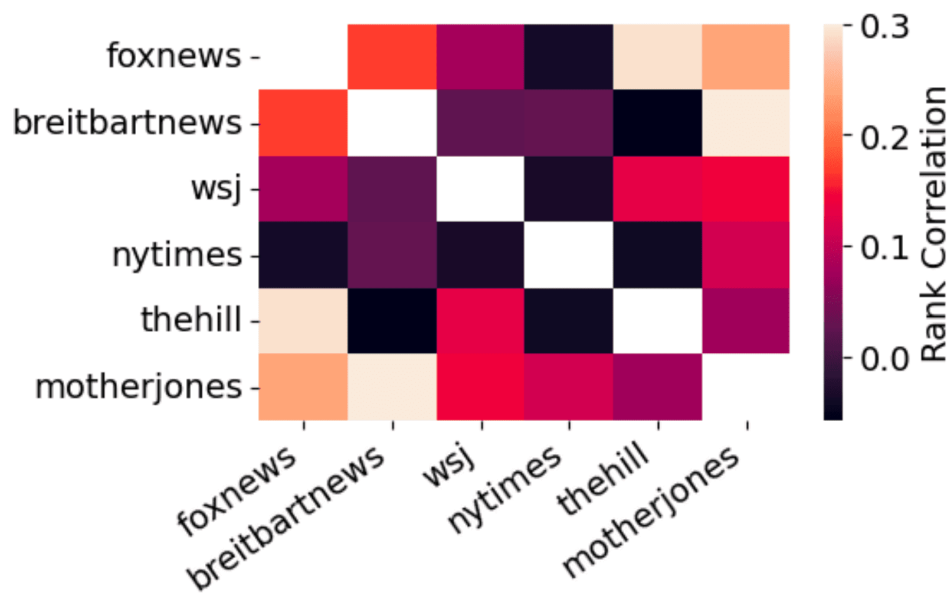
To evaluate the practical utility of our models, we design two downstream tasks: (1) analyzing newsworthiness agreement between publishers, and (2) using newsworthiness models to rank corpora of interest to journalists.

### 5.1. Task 1: Newsworthiness Agreement Between Publishers


In this task, we aim to rank-order lists of news items drawn from a larger pool of articles to calculate the agreement rates for newsworthiness decisions between different news outlets. Previous research has observed surprising overlaps in sentiment and preferences between right-leaning and left-leaning outlets<sup>[26]</sup>, and we wish to quantitatively test this phenomenon using our preference models.

We selected 9 of the 31 outlets for which we trained preference models in the previous section. From each outlet, we sampled 1,000 articles, matching on variables such as topic, length, publication date, and other potential confounders. These 9 outlets were chosen because they represent a range of political viewpoints.

For each model  $n_{o_i}$  (corresponding to outlet  $o_i$ ), we used it to sort lists of 1,000 articles  $\{a_1, a_2, \dots, a_{1000}\}_{j=1}^9$  from outlets  $\{o\}_{j=1}^9$ . In other words, the output of applying model  $n_{o_i}$  to the article list from outlet  $o_j$  is a fully sorted list  $n_{o_i}(A_j)$ . We used the size  $\times$  position model for this experiment, as performance was similar to the size-only model, and we believed that the multivariable models capture more newsworthiness information than the single-variable models.



**Figure 2a.** Kendall's  $\tau$  correlation between the newsworthiness preferences expressed by preference models trained on different news outlets.



**Figure 2b.** Cosine distance of average SBERT similarity between articles sampled from each outlet.

**Figure 2.** Comparison of Kendall’s  $\tau$  rank correlation (on newsworthiness judgements) and SBERT cosine similarity (on articles) across news outlets.

We calculated Kendall’s  $\tau$ , a correlation measure for ordinal data, between each pair of sorted lists  $(n_{o_i}(A_k), n_{o_j}(A_k))$  for all  $i, j, k$ , and averaged the correlations across  $j$ . The resulting correlation matrix is displayed in Figure 2a. Some surprising insights emerge from this analysis. Notably, Breitbart, a right-leaning outlet, and Mother Jones, a left-leaning outlet, have one of the highest rates of agreement.

To establish a baseline and ensure we are not merely capturing topic overlap (despite matching on topics), we conducted a simple SBERT embedding experiment<sup>[27]</sup>. We sampled a set of 100 articles per outlet, generated embeddings using SBERT, and averaged these embeddings to create single outlet-level embeddings, as shown in Figure 2b. These embedding-level similarities align more closely with topical overlaps, indicating distinct right-wing and left-wing clusters with some overlap in between.

Taken together, these results suggest that newsworthiness preference is a compelling and orthogonal variable for study beyond topical similarity.

## 5.2. Task 2: Surfacing Potentially Newsworthy Leads

In this task, we explore how well these newsworthiness judgments transfer outside of the news domain. In this task, we build on the work of<sup>[5]</sup>. The authors introduced the task of *newsworthiness prediction* as a detection and alerting system for journalists: utilizing a list of San Francisco Board of Supervisors’ policies (a typical source of stories for journalists), they attempted to detect which policies were more newsworthy in order to alert journalists.

Here, we suspect that editorial cues from different homepages will help us surface especially newsworthy content based on the preferences of each outlet. We applied the models from each outlet to sort the list of Board of Supervisors’ policies. Then, we selected the top 10 items from the ordered lists  $n_{o_i}$  and used a large language model (LLM) to summarize the key points raised in each policy.<sup>9</sup>

The LLM’s summarization results and examples are shown in Table 5. Themes emerge, with subject-specific outlets like *The Weather Channel* highlighting policies related to environmental issues. We presented these results to a group of journalists, and 81% of respondents indicated they were impressed and would consider using such a system in their workflow. These findings demonstrate the



potential of our models to assist journalists in identifying newsworthy leads from large corpora of documents, thereby supporting investigative journalism and timely reporting.

Outlet	Top Policies LLM Summaries	Examples of Policies
Weather Channel	Environmental Policies, Public Health and Emergency Response, Infrastructure and Development	Reducing nutrient pollution from wastewater; Accepting grants for forensic science improvements
Daily Climate	Environmental and Energy Policies, Urban Planning and Development	Agreement with North Star Solar; Building code enforcement
Fox News	Community and Public Safety Policy, Education and Social Policy, Fiscal and Economic Policy	Appointment of individuals to advisory committees; Appropriating funds for San Francisco Unified School District; Developing materials on domestic violence
Mother Jones	Social Policies, Environmental and Health Policies	Sanctuary City Protection; Urging Pardons; Edible Food Recovery and Organic Waste Collection
Ars Technica	Infrastructure Policies	System Impact Mitigation Agreement; 6th St. Substation
NYTimes	Social & Cultural Awareness Policies, Labor & Employment, Economic, Housing policies	Commemorative and Awareness Events; Labor Dispute Hearings; Affordable Housing Loans
WSJ	Economic and Infrastructure Policies, Governance and Legislative Policies	Contract modifications; Bond sales; Ground lease agreements; Charter amendments concerning commissions and departments related to aging and adult services

**Table 5.** Summaries of the top 10 most newsworthy policies published by the San Francisco Board of Supervisors, as ranked by models trained on 7 different homepages.

## 6. Discussion

Our demonstrations show two core findings: first, editorial priorities and decision-making can be inferred simply by examining the layout decisions made on homepages. This decision-making is

distinct from simple topic preferences, as we show. In fact, commonalities about decision-making can be observed between outlets that appear distinct topically. Second, newsworthiness judgements have potential to be used in tools for reporting.

Stepping back, these results indicate that homepage editorial cues provide an interesting, novel angle for news analysis, as well as a tantalizing direction in newsworthiness detection<sup>[5][28]</sup>. Both of these applications are premised by the assumption that editorial cues learned from one outlet’s homepage can be transferred to other domains, be it *another* outlet’s articles, or non-news content. This is an important assumption: the intuitive findings that we have made in our demonstrations provide some degree of proof that this transfer is robust.

We experimented with different ways of making this transfer even more robust. We attempted to train *additional* models to serve as in-domain and out-of-domain classifiers, and then multiplied the probabilities:  $\hat{P}_o(a_1 > a_2) = p_o(in\_domain|a_1, a_2) \times p_o(a_1 > a_2)$ , where for outlet  $o$ ,  $p_o(in\_domain|a_1, a_2) = 1$  if  $a_1$  belongs to  $o$  while  $a_2$  does not. However, our results in this direction were not more interpretable than the results we reported. Ultimately, without any “gold truth” about how an editor from outlet  $o$  would rank an arbitrary list of strings, we will not have a conclusive measurement about our ability to replicate these measurements. In order to fully validate our observations, this appears necessary. Our results have to be taken with some important further caveats. While some newsworthiness models were impressively well-performing, many were not. Further exploration is needed to determine the causes. Additionally, despite the presence of non-English homepages in our dataset, we only tested with U.S.-based websites. We look forward to continuing to expand this work to address these concerns.

With these caveats in mind, we hesitate to draw firm conclusions. However, we feel the results of our modeling are encouraging enough to continue to address these concerns and run further manual experiments with editors. We imagine a future where editorial preferences made by professional editors can be used to routinely study importance and organization of content. We imagine insights being applied more broadly to build tools for journalists, improve webpage layouts that are currently automated, and even understand more fundamental components of the human psyche.

## 7. Related Work

Understanding how information is prioritized and presented in news media has long been a subject of scholarly interest. This section situates our work within the broader literature on information

prioritization, visual cues in editorial decision-making, layout parsing, modeling editorial judgments, and data-driven studies of news content and bias.

### *Information Prioritization and Newsworthiness*

The concept of *newsworthiness* is central to journalism studies and media sociology. Classic works by<sup>[29]</sup> introduced a set of news values that determine the selection and presentation of news stories. These news values have been revisited and updated by scholars such as<sup>[30]</sup> and<sup>[31]</sup>, who identified factors like relevance, timeliness, and unexpectedness as key determinants of newsworthiness. Prior research has explored the cognitive and organizational processes behind news selection. <sup>[32]</sup> examined the gatekeeping role of editors and journalists in filtering news content. <sup>[33]</sup> discussed how media organizations' structures influence news production and prioritization.

Our work contributes to this literature by providing a computational approach to inferring newsworthiness judgments from homepage layouts, offering a large-scale empirical perspective on editorial prioritization decisions.

### *Visual Cues and Editorial Decision-Making*

Visual presentation plays a crucial role in shaping readers' perceptions of news importance. Studies have shown that elements such as headline size, article position, and the use of images significantly affect reader attention and recall <sup>[34]</sup><sup>[35]</sup>.

Eye-tracking research has provided insights into how users interact with news websites. <sup>[23]</sup> found that users' viewing patterns are influenced by the layout and design of webpages, with a tendency to focus on content positioned at the top and left areas of the screen. Similarly, <sup>[10]</sup> demonstrated that visual cues guide readers' attention and are integral to the perceived relevance of news stories. Our study builds on these findings by quantitatively modeling how spatial attributes of articles on homepages reflect editorial judgments about their significance.

### *Modeling Editorial Decisions*

Computational modeling of editorial decisions and news selection processes has been explored in prior research. <sup>[36]</sup> proposed methods to learn news selection patterns from data, aiming to predict which news stories editors might choose to publish. Additionally, <sup>[28]</sup> investigated computational

tools to assist journalists in identifying newsworthy data points within large datasets, emphasizing the role of algorithms in supporting editorial judgment.

Our approach differs by focusing on inferring editorial prioritization from the spatial arrangement of articles on homepages, rather than predicting content selection from textual features alone. By modeling pairwise preferences based on layout attributes, we provide a novel perspective on editorial decision-making.

### *Data-Driven Studies of News Content and Bias*

Large-scale analyses of news content have provided insights into media bias, framing, and agenda-setting.<sup>[26]</sup> examined the factors influencing media slant in U.S. daily newspapers, using textual analysis to measure ideological positioning. Multiple research projects have aggregating and analyzing vast amounts of news content, enabling researchers to explore patterns in news coverage and discourse<sup>[37][38][39][40][41]</sup>.

Our dataset, NewsHomepages, contributes to this line of research by offering a rich source of data on how news organizations present and prioritize information on their homepages, complementing textual analyses with spatial and visual dimensions.

### *Applications in Journalism and News Recommendation*

The intersection of computational models and journalism has led to the development of tools for news recommendation and content personalization.<sup>[42]</sup> explored context-aware news recommendation systems that adapt to users' interests and reading behaviors. Moreover, computational approaches have been proposed to assist journalists in investigative work.<sup>[5][43][44]</sup> developed methods to detect newsworthy events in local government proceedings, identify sources, and structure news articles.

Our work extends these applications by demonstrating that models trained on editorial cues from news homepages can be applied to rank a variety of corpora, possibly aiding in the design of future recommendation systems.

### *Visual Salience and Information Design*

The principles of information design and visual salience are relevant to our study.<sup>[20]</sup> emphasized the impact of visual arrangement in effectively conveying information. In the context of web design, <sup>[45]</sup> discussed how layout aesthetics affect user engagement and perception.

Our findings align with these principles, highlighting how the spatial arrangement of content on news homepages serves as an implicit signal of importance, guiding reader attention and shaping information consumption.

## 8. Conclusion

In this work, we introduced NewsHomepages, a comprehensive dataset capturing editorial prioritization decisions through the lens of homepage layouts across thousands of news organizations. By applying weakly-supervised models, we demonstrated the ability to robustly infer editorial judgments through pairwise comparisons of article positioning.

## 9. Limitations

This work, while advancing the study of editorial prioritization on homepages, comes with several limitations. First, the dataset, although large and diverse, predominantly focuses on English-language news outlets from the U.S., which may limit the generalizability of our models to international or non-English outlets. Despite the inclusion of some non-U.S. and non-English homepages, the models have not been explicitly evaluated on a broader range of languages or cultural contexts. This focus may overlook regional editorial conventions and biases that differ significantly from the U.S. context.

Another limitation is that the study focuses primarily on visual cues of newsworthiness, such as size, position, and graphical elements. While these cues are significant, they are not the only factors that influence editorial decisions. The models do not account for less visible but equally critical considerations, such as journalistic ethics, editorial mandates, or audience engagement metrics, which may influence homepage layouts but remain unquantified in this dataset.

Additionally, the weakly-supervised learning methods employed for layout parsing may struggle with more complex or irregular homepage designs. As described, the model had trouble generalizing to homepages with obscure HTML structures, leading to imperfect bounding box detections in some cases. This could result in misinterpretations of editorial significance, especially for websites with non-traditional or highly dynamic layouts.

Lastly, the results may have some bias due to the reliance on pairwise article comparisons. This method, while efficient, reduces the complexity of editorial decision-making into binary

relationships, potentially overlooking more nuanced or multifaceted prioritization strategies that editors use in practice.

### *9.1. Computational Budget*

The computational resources required for this project were substantial. The model training involved  $4 \times A40$  GPUs for initial phases and  $16 \times A100$  GPUs for more extensive model fine-tuning and deployment. Training the custom Detectron2 model alone took 24 hours, with additional time required for fine-tuning and model testing across multiple outlets. While the experiments could be completed within this budget, more extensive experiments across all 3,000 outlets would require significantly more resources, especially when scaling to different languages and regional variations.

### *9.2. Use of Annotators*

The dataset used in this work was primarily compiled automatically through web scraping, with minimal manual annotation. However, annotators were employed to manually label URLs to distinguish between news articles and non-news articles during the preprocessing phase. A set of 2,000 URLs was manually labeled, contributing to a more refined and accurate dataset. The authors performed this task themselves. However, beyond this, no human annotators were used to manually assess newsworthiness, as the models relied on position and size as proxies for editorial decisions.

## **Appendix A. Dataset Details**

In this section, we present more detailed dataset statistics. In Table 6, we show the different languages collected in our corpora and in Table 8 we show

'ainonline, airwaysmagazine, arstechnica, bleacherreport, breitbartnews, chessbase, cnet, denverpost, foxnews, jaxdotcom, jessicavalenti, jezebel, motherjones, newsobserver, nytimes, phoenixluc, rollcall, seattletimes, sfchronicle, sinow, slate, starteleggram, studyfindsorg, thealligator, theathletic, thedailyclimate, thehill, weatherchannel, wired, wsj, yaledailynews'

Language	Count
English	975
Spanish, Castilian	44
Portuguese	36
Nepali	24
French	21
German	10
Japanese	9
Norwegian	8
Hindi	7
Hebrew	7
Russian	7
Italian	5
Ukrainian	5
Chinese	3
Afrikaans	3
Zulu	2
Xhosa	1

**Table 6.** Our corpus comprises homepages from 18 different languages. We assign each news outlet to the language of the majority of its articles' languages (e.g. the New York Times sometimes publishes Spanish-language articles, but is predominantly an English-language newspaper).

Outlet	Domain		Position x Size	
	Train	Test	Train	Test
ainonline	2159	540	3844	962
airwaysmagazine	1233	309	1669	418
arstechnica	9349	2338	17883	4471
bleacherreport	3849	963	6689	1673
breitbartnews	7824	1957	15199	3800
chessbase	2094	524	3151	788
cnet	3769	943	6521	1631
denverpost	18802	4701	36607	9152
foxnews	62170	15543	125096	31274
jaxdotcom	4100	1026	7206	1802
jessicavalenti	455	114	512	129
jezebel	6270	1568	10956	2740
motherjones	2443	611	3572	893
newsobserver	7538	1885	14078	3520
nytimes	38432	9608	74857	18715
phoenixluc	209	53	305	77
rollcall	6572	1643	12436	3109
seattletimes	20942	5236	40882	10221
sfchronicle	5600	1401	10952	2739
sinow	176	44	315	79
slate	23527	5882	45470	11368
startelegram	11964	2992	22932	5734
studyfindsorg	450	113	403	101
thealligator	3046	762	5432	1359



Outlet	Domain		Position x Size	
	Train	Test	Train	Test
theathletic	22722	5681	44463	11116
thedailyclimate	7211	1803	13688	3423
thehill	30800	7700	59965	14992
weatherchannel	4551	1138	8094	2024
wired	2058	515	3196	800
wsj	15569	3893	30496	7625
yaledailynews	1378	345	2527	632

**Table 7.** Size of training and test sets, in terms of # of pairs, used in our experiments.

Country	Count
United States	892
Brazil	37
United Kingdom	32
Nepal	24
Canada	20
South Africa	18
France	17
Spain	13
Mexico	13
India	10
Japan	9
Argentina	9
Israel	9
Germany	9
Russia	8
Norway	8
Ukraine	6
Ireland	6
Italy	5
New Zealand	4
Austria	3
Taiwan	3
Colombia	2
Australia	2
Uruguay	1

Country	Count
Qatar	1
Belgium	1
Latvia	1
Bosnia and Herzegovina	1
Georgia	1
El Salvador	1
Lebanon	1

**Table 8.** Countries of origin for the homepages we collect, based on where the organization is based.


## Appendix B. News Homepage Layouts

In Figure 3, we show several areas of a homepage where editorial policies are likely to be unclear and challenging to model. In the top section, a “Breaking News” feed shows articles shortly after they are published. They usually do not stay long in these positions<sup>[46]</sup>, so there is high variability in this section. In the middle section, a “Section Fronts” show top articles in each section, each combining the different priorities of the desks<sup>[19]</sup>. Finally, in this reporter’s experience, the bottom of a homepage was affectionately called the “Gutter”. However, it is more commonly referred to as a “footer”<sup>10</sup>



**Figure 3.** We show three sections of a sample homepage (from CBS News) where editorial decisions for different reasons. We highlight the “Breaking News” Section, “Section Fronts” and “The Footer”.

We can extend this analysis by visualizing the flow of articles on a homepage over time. We show in Figure 5 where articles tend to get added and deleted overall, as well as where they stay the longest. In Figure 6, we show how articles shift frequently around a homepage.

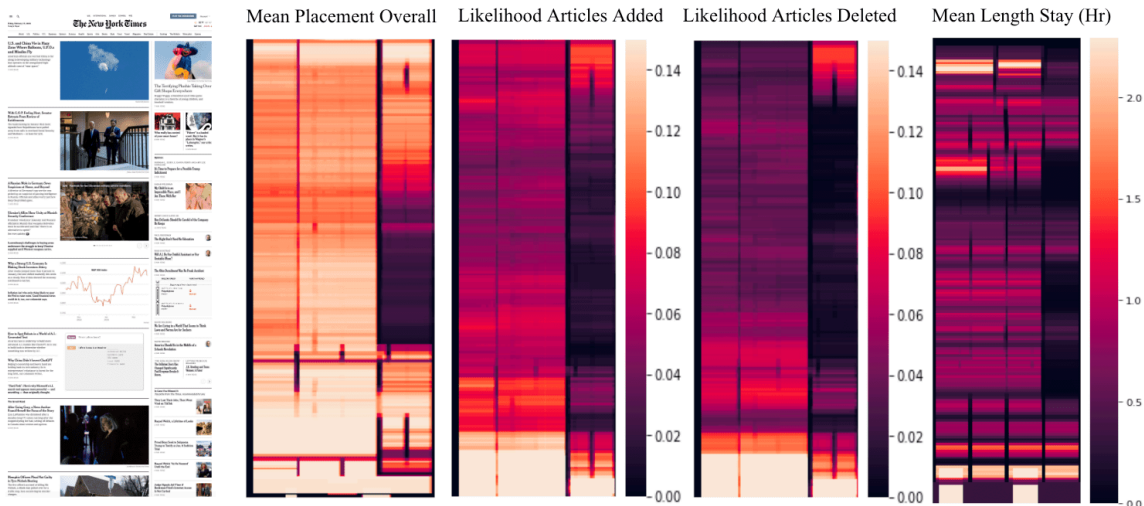


(a) Our deterministic algorithm starts at all  $\langle a \rangle$  nodes and recursively traverses up the DOM to find maximal subtrees with one  $\langle a \rangle$ . Green nodes shown are article bounding boxes.

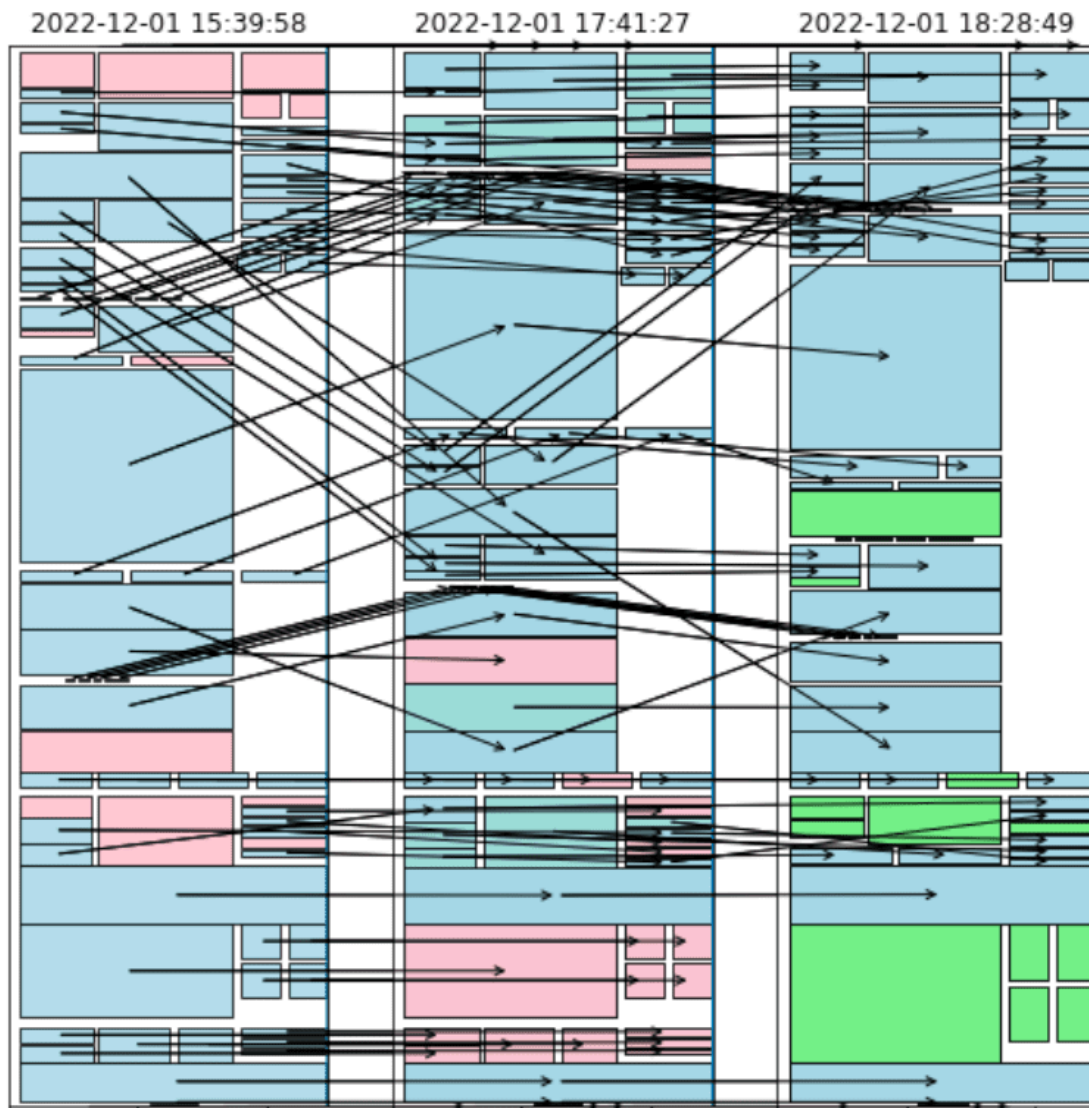


(b) Failure cases (missing text area) with the deterministic algorithm.

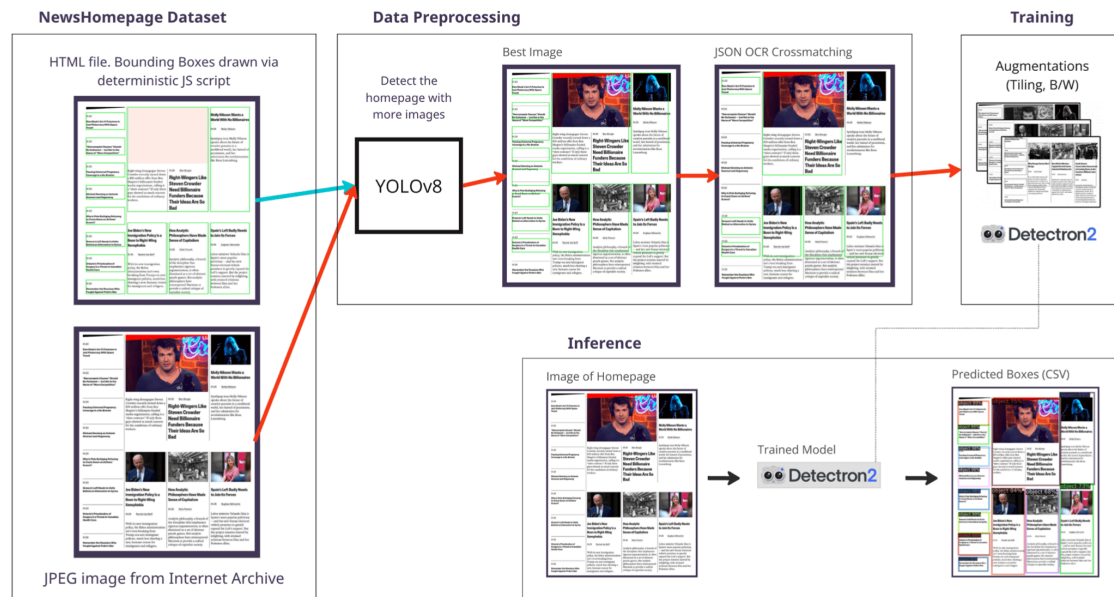
**Figure 4.** Illustration of our deterministic bootstrapping algorithm and a failure case. Here, when non-article links exist, we misunderstand the full area of an article, excluding the text below.




**Figure 5.** Different analyses we run on bounding boxes across time: average locations of bounding boxes on a homepage, locations where articles are added first, locations where they are removed, and the average time articles in various locations spend.



**Figure 6.** With our suite of tools for parsing homepages, we can examine on a granular level the movement of an article across the homepage.



**Figure 7.** This diagram is an overview of the data preparation and training of the Detectron2 model for predicting bounding box on websites.



**Figure 8.** When sorting our sources to determine the ones most difficult for the DOM-Tree algorithm, we define the Average Similarity score to be a general measure as to how well the bounding box's text match the article's JSON file containing text/link pairs. High similarity score means high bounding box accuracy, and vice versa.



## Statements and Declarations

**Co-first author contributions:** Ben Welsh founded the News-Homepages Internet Archive Project, built and maintained the pipelines to perform data collection over a period of 8 years. Alexander Spangher was the primary writer, he was the advisor for Arda and Michael, and he performed the analysis in Section 5. Arda Kaz did the modeling in Section 4. **Secondary author contributions:** Michael Vu did most of Section 3. Naitian performed initial exploration in Section 5 and editing.

## Footnotes

<sup>1</sup> Terms like “above the fold” emerged to signal story-importance (i.e. the story is above the point at which the newspaper folds, so it is seen on newsstands).

<sup>2</sup> For more information on how to contribute, please see: <https://github.com/palewire/news-homepages>. For all code and data associated with this project, see <https://github.com/alex2awesome/homepage-newsworthiness-with-internet-archive>.

<sup>3</sup> <https://github.com/gildas-lormeau/SingleFile>, incidentally the same software that Zotero uses. In initial experimentation, we observed that capturing complete, compressed HTML snapshots was far more robust than capturing assets.

<sup>4</sup> <https://archive.org/details/news-homepages>.

<sup>5</sup> Likely due to errors in HTML extraction or dead links.

<sup>6</sup> Figure 8, in the Appendix, demonstrates the resulting histogram for the distribution of OCR-match scores across our dataset.

<sup>7</sup> In Appendix B.

<sup>8</sup> As classified by MediaBiasFactCheck.com

<sup>9</sup> We used GPT-4 for this experiment.

<sup>10</sup> <https://alyamanalhayekdesign.com/blog/the-parts-of-a-webpage-a-complete-list/>

## References

- <sup>1</sup> <sup>^</sup>Miller GA (1956). "The magical number seven, plus or minus two: Some limits on our capacity for processing information." *Psychological review*. 63 (2): 81.

2. <sup>△</sup>Rosenfeld L. *Information architecture for the world wide web*. O'Reilly Media; 2002.
3. <sup>△</sup>Hays S (2018). "An Analysis of Design Components of Award-winning Newspaper Pages". *Elon Journal of Undergraduate Research in Communications*. 9 (2): 44–63.
4. <sup>△</sup>Boukes M, Jones NP, Vliegenthart R (2022). "Newsworthiness and story prominence: How the presence of news factors relates to upfront position and length of news stories". *Journalism*. 23(1):98–116.
5. <sup>△</sup><sub>a, b, c</sub> <sup>△</sup>Spangher A, Ferrara E, Welsh B, Peng N, Tumgoren S, May J (2023). "Tracking the Newsworthiness of Public Documents". *arXiv preprint arXiv:2311.09734*.
6. <sup>△</sup>Barnhurst KG, Nerone J. *The Form of News: A History*. Guilford Press; 2001.
7. <sup>△</sup>Usher N. *Making news at the New York times*. University of Michigan Press; 2014.
8. <sup>△</sup>Sullivan M (2016). "The End of the Page One Meeting: Making Way for the Reader in Choosing the News". *The New York Times*. Available from: <https://publiceditor.blogs.nytimes.com/2016/03/16/the-end-of-the-page-one-meeting-making-way-for-the-reader-in-choosing-the-news/>. Accessed 2023 Oct 6.
9. <sup>△</sup>Nielsen J. *F-Shaped Pattern For Reading Web Content*. 2006. Available from: <https://www.nngroup.com/articles/f-shaped-pattern-reading-web-content/>. Accessed: 2023-10-06.
10. <sup>△</sup><sub>a, b</sub> <sup>△</sup>Bucher HJ, Schumacher P (2006). "The relevance of attention for selecting news content. An eye-tracking study on attention patterns in the reception of print and online media." *Walter de Gruyter*. 2006.
11. <sup>△</sup>García MR. *Contemporary Newspaper Design: Shaping the News in the Digital Age*. Prentice Hall; 1987.
12. <sup>△</sup>Zillmann D, Knobloch S, Yu Z (2001). "Effects of photographs on the selective reading of news reports". *Media Psychology*. 3 (4): 301–324.
13. <sup>△</sup>Shen Z, Zhang R, Dell M, Moser BC, Carlson J, Li W (2021). "LayoutParser: A Unified Toolkit for Deep Learning Based Document Image Analysis". In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM. pp. 4528–4538.
14. <sup>△</sup>Li M, Cui L, Huang S, Wei F, Zhou M (2020). "DocBank: A Benchmark Dataset for Document Layout Analysis". *Proceedings of the 28th International Conference on Computational Linguistics*. pages 949–960.
15. <sup>△</sup>Zhong X, Tang J, Yepes AJ. *PubLayNet: largest dataset ever for document layout analysis*. In: *2019 International Conference on Document Analysis and Recognition (ICDAR)*. IEEE; 2019. p. 1015–1022.
16. <sup>△</sup><sub>a, b</sub> <sup>△</sup>Amini MR, Feofanov V, Pauletto L, Hadjadj L, Devijver E, Maximov Y (2022). "Self-training: A survey". *arXiv preprint arXiv:2202.12040*.

17. <sup>△</sup>Redmon J, Farhadi A (2018). "YOLOv3: An Incremental Improvement". arXiv.
18. <sup>△</sup>Wu Y, Kirillov A, Massa F, Lo WY, Girshick R (2019). Detectron2. Available from: <https://github.com/facebookresearch/detectron2>.
19. <sup>△</sup>Angell C. Metrics at work: Journalism and the contested meaning of algorithms. Princeton University Press, Princeton, NJ; 2020.
20. <sup>△</sup>Tufte ER. Envisioning Information. Graphics Press; 1990.
21. <sup>△</sup>Salganik MJ, Dodds PS, Watts DJ (2006). "Experimental study of inequality and unpredictability in an artificial cultural market". *Science*. 311 (5762): 854–856.
22. <sup>△</sup>Lambert A, Brock J (2005). "Layout complexity and visitors' attention on web pages: An eye-tracking study". *Journal of Digital Information*. 6 (2).
23. <sup>△</sup>Nielsen J, Pernice K (2009). Eyetracking Web Usability. New Riders.
24. <sup>△</sup>Spangler A (2015). "Building the next New York Times recommendation engine". *The New York Times*. pages 08–26.
25. <sup>△</sup>Mangrulkar S, Gugger S, Debut L, Belkada Y, Paul S, Bossan B (2022). "PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods". Available from: <https://github.com/huggingface/peft>.
26. <sup>△</sup>Gentzkow M, Shapiro JM (2010). "What drives media slant? Evidence from US daily newspapers". *Econometrica*. 78 (1): 35–71.
27. <sup>△</sup>Reimers N, Gurevych I (2019). "Sentence-BERT: Sentence embeddings using Siamese BERT-networks". In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 3982–3992.
28. <sup>△</sup>Diakopoulos N, Naaman M, Kivran-Swaine F (2010). "Diamonds in the rough: Social media visual analytics for journalistic inquiry." In: *2010 IEEE Symposium on Visual Analytics Science and Technology*. IEEE. pp. 115–122.
29. <sup>△</sup>Galtung J, Ruge MH (1965). "The structure of foreign news". *Journal of Peace Research*. 2 (1): 64–90.
30. <sup>△</sup>Harcup T, O'Neill D (2001). "What is news? Galtung and Ruge revisited". *Journalism Studies*. 2 (2): 261–280.
31. <sup>△</sup>Harcup T, O'Neill D (2017). "What is news? News values revisited (again)". *Journalism Studies*. 18 (12): 1470–1488.
32. <sup>△</sup>Shoemaker PJ. Gatekeeping. Sage Publications; 1991.

33. <sup>△</sup>Herman ES, Chomsky N. "Manufacturing consent." In: *Power and Inequality*. Routledge; 2021. p. 198–206.
34. <sup>△</sup>Brooks BS, Pinson JL. *The Art of Editing: In the Age of Convergence*. Routledge; 2022.
35. <sup>△</sup>Nass C, Mason L (1990). "On the study of technology and task: A variable-based approach". *Organizations and communication technology*. 46: 67.
36. <sup>△</sup>Arya C, Dwivedi SK. News web page classification using url content and structure attributes. In: 2016 2nd International Conference on Next Generation Computing Technologies (NGCT). IEEE; 2016. p. 317–322.
37. <sup>△</sup>Roberts H, Bhargava R, Valiukas L, Jen D, Malik MM, Bishop CS, Ndulue EB, Dave A, Clark J, Etling B, et al. Media cloud: Massive open source collection of global news on the open web. In: *Proceedings of the International AAAI Conference on Web and Social Media*. 2021; 15:1034–1045.
38. <sup>△</sup>Misra R. News Category Dataset. arXiv preprint arXiv:2209.11429. 2022.
39. <sup>△</sup>Silcock E, Arora A, D'Amico-Wong L, Dell M (2024). "Newswire: A Large-Scale Structured Database of a Century of Historical News". arXiv preprint arXiv:2406.09490.
40. <sup>△</sup>Leetaru K, Schrodtt PA (2013). "Gdelt: Global data on events, location, and tone, 1979–2012". In: *ISA annual convention*, volume 2, pages 1–49. Citeseer.
41. <sup>△</sup>Spangher A, Ren X, May J, Peng N (2022). "NewsEdits: A news article revision dataset and a novel document-level reasoning challenge". *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Seattle, United States: Association for Computational Linguistics. pp. 127–157. doi:[10.18653/v1/2022.naacl-main.10](https://doi.org/10.18653/v1/2022.naacl-main.10).
42. <sup>△</sup>TalebiFard P, Leung VC (2014). "Context-aware dissemination of information and services in heterogeneous network environments". *Journal of Ambient Intelligence and Humanized Computing*. 5: 775–787.
43. <sup>△</sup>Spangher A, Peng N, Ferrara E, May J (2023). "Identifying informational sources in news articles". *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore: Association for Computational Linguistics. pp. 3626–3639. doi:[10.18653/v1/2023.emnlp-main.221](https://doi.org/10.18653/v1/2023.emnlp-main.221). Available from: <https://aclanthology.org/2023.emnlp-main.221>.
44. <sup>△</sup>Spangher A, Ming Y, Hua X, Peng N (2022). "Sequentially controlled text generation". *Findings of the Association for Computational Linguistics: EMNLP 2022*. pages 6848–6866. doi:[10.18653/v1/2022.findings-emnlp.509](https://doi.org/10.18653/v1/2022.findings-emnlp.509).

45. <sup>^</sup>Jiang Z, Wang W, Tan BC, Yu J (2016). "The determinants and impacts of aesthetics in users' first interaction with websites". *Journal of Management Information Systems*. 33 (1): 229–259.
46. <sup>^</sup>Costanza-Chock S, Rey-Mazon P (2016). "PageOneX: New approaches to newspaper front page analysis". *International Journal of Communication*. 10: 28.

## Declarations

**Funding:** No specific funding was received for this work.

**Potential competing interests:** No potential competing interests to declare.