

Paper Skygest: Personalized Academic Recommendations on Bluesky

Sophie Greenwood
sjgreenwood@cs.cornell.edu
Cornell Tech
New York, New York, USA

Nikhil Garg
ngarg@cornell.edu
Cornell Tech
New York, New York, USA

ABSTRACT

We build, deploy, and evaluate PAPER SKYGEST, a custom personalized social feed for scientific content posted by a user’s network on Bluesky and the ATProtocol. We leverage a new capability on emerging decentralized social media platforms: the ability for anyone to build and deploy feeds for other users, to use just as they would a native platform-built feed. To our knowledge, PAPER SKYGEST is the first and largest such continuously deployed personalized social media feed by academics, with over 35,000 weekly uses by almost 1200 daily active users, all organically acquired. First, we quantitatively and qualitatively evaluate PAPER SKYGEST usage, showing that it has sustained usage and satisfies users; we further show adoption of PAPER SKYGEST increases a user’s interactions with posts about research, and how interaction rates change as a function of post order. Second, we share our full code and describe our system architecture, to support other academics in building and deploying such feeds sustainably. Third, we overview the potential of custom feeds such as PAPER SKYGEST for studying algorithm designs, building for user agency, and running recommender system experiments with organic users without partnering with a centralized platform.

1 INTRODUCTION

Social media recommendation feeds – on Facebook, Twitter/X, Bluesky, Reddit, TikTok, etc – mediate individual and societal consumption and discussion of news stories, scientific content, and entertainment [21, 77, 87]. These feeds are largely *centrally designed and deployed*, with users given the choice between only a few “algorithmic” feeds developed by the central platform and a “non-algorithmic” version consisting of posts by accounts the user follows, in reverse chronological order.

This centralization limits academic study of algorithm design and causal effects of feeds on such platforms [53, 62, 68]. (1) First, due to bureaucratic and engineering constraints: researchers must either (a) partner with the central platform to study how algorithm design affects organic users [29, 30, 37, 44, 52, 54, 63, 75], or (b) externally run controlled lab experiments on simulated platforms or develop a custom layer on top of a platform (such as through browser extensions) [12, 16, 18, 40, 43, 58, 68, 69, 73, 86]. Platform partnership opportunities are rare,¹ and experimental designs are constrained by the platform’s interests. Lab experiments require recruiting a set of participants (often paid), and the design space for browser extensions is often limited to *re-ranking* content surfaced by the platform’s feeds [68, 69].² Maintaining continuous usage

for external platforms is a further challenge. (2) Second, and more fundamentally, a common design characteristic is that a single feed algorithm often surfaces *all* types of content – whether political, research, entertainment, or social. Thus, studying algorithmic amplification – either with a platform or using a browser extension – is often limited to topics such as news consumption and polarization, which are assumed to affect a large fraction of the overall user base.

This work leverages a new opportunity for social media research that allows academic researchers to flexibly control user recommendations in a real platform environment: *custom feeds* on Bluesky [50]. Bluesky is an emerging social media platform, on the ATProtocol, where users publish and interact with short-form text posts (like Twitter/X), with about 40 million total users and over 1 million daily active users as of October 2025. Bluesky allows anyone to host arbitrary recommendation algorithms, that other users can (organically) choose to use – fulfilling a desire for such decentralized “middleware” [23]. As shown in Figure 1a, custom feeds pinned by a user are shown alongside platform-provided feeds. For example, there are popular feeds that are dedicated to trending news, various academic communities (machine learning/AI, Economics), and “Quiet Posters,” which prioritizes users who do not post much. Tools such as Graze.social allow anyone to build feeds without writing code and host thousands of custom feeds, many with tens of thousands of users.

We leverage this opportunity, and make three contributions. (1) First, we build, deploy, and evaluate PAPER SKYGEST, a custom *personalized* feed for scientific content: it shows posts about research publications by those in the user’s following network. To our knowledge, PAPER SKYGEST is the first (and perhaps the only, to date) personalized social media feed deployed at scale for organic usage on a major social media platform either by academics or for scientific content.³ PAPER SKYGEST serves scientists and others who discuss scientific content on Bluesky. Social media, especially “academic Twitter,” has long driven consumption of academic content. Especially since Fall 2024, many academics have migrated from Twitter/X to Bluesky [78], although adoption is heterogeneous between research communities [72]. Altmetric, which tracks social media mentions of published papers, estimated that for most of March 2025 Bluesky hosted as many or more daily links to published academic research as did Twitter/X [84].

¹As González-Bailón et al. [25] state, “There is virtually no research comparing what people could potentially see within platforms and what they actually see [with the exception of Bakshy et al. [4]].”


²As Piccardi et al. [68] explain regarding browser experiments, “without access to the complete platform inventory, up-ranking content in the feed presents more challenges

than down-ranking content. The content returned by the server is obtained by scoring the inventory, which includes the entire candidate set of potential posts available for display. Without the full set or a meaningful approximation, the intervention can operate only on the pre-selected content delivered to the user, representing the top items the curation algorithm has already prefiltered.”

³Feeds such as the Science feed or various AI/ML feeds show the same content to everyone. In general, building personalized feeds is much more challenging than non-personalized ones, which can be built using “no code” solutions such as Graze.



(a) PAPER SKYGEST as it appears to a user in the Bluesky interface.



(b) PAPER SKYGEST Infrastructure.

Figure 1: (a) A user’s perspective of using PAPER SKYGEST: the user switches to the feed and sees a list of posts about research papers (including the broader thread) by people in their following network. (b) A diagram of the PAPER SKYGEST backend, which is comprised of three main components: (1) the *feed generator endpoint*, which serves cached recommendations to users, (2) the *firehose*, which listens for and stores paper posts, and (3) the *recommendation generation* module, which precomputes the paper posts to show to each PAPER SKYGEST user at regular intervals and caches them.

PAPER SKYGEST has continuously run since its public launch in March 2025 and has garnered substantial organic user adoption, with over 35,000 weekly uses by almost 1,200 daily active users, totaling over 1 million uses to date. We show qualitatively and quantitatively that users find PAPER SKYGEST useful and continuously return to the feed; we further use PAPER SKYGEST data to show how users consume and interact with scientific content, and that using PAPER SKYGEST shifts consumption patterns toward scientific content.

(2) To enable other academics to build, deploy, and evaluate such feeds for themselves, we describe the engineering architecture behind PAPER SKYGEST. We describe engineering lessons learned, especially for other academics aiming to build such feeds for research. Our code is available at <https://github.com/Skygest/PaperSkygest>.

(3) We provide a vision for how academics can use new capabilities provided by the ATProtocol and Bluesky to run large-scale randomized field experiments on algorithmic design with organic users, without partnering with a platform. This new capability enables deployments and evaluations of new approaches to support user agency, discovery of niche topics, and feed ranking beyond engagement optimization.

This paper is organized as follows. Related work is in Section 2. In Section 3, we overview PAPER SKYGEST and present quantitative usage metrics and qualitative feedback; we further leverage PAPER SKYGEST data to show how users view, consume, and share research on Bluesky. Section 4 contains a description of our system architecture. In Section 5, we describe how the ATProtocol, and PAPER SKYGEST in particular, enables new forms of social media algorithm design and evaluations. We conclude in Section 6.

2 RELATED WORK

Deployed Recommenders from Academia. Industry building and deploying recommender systems is common, with many publications on experimental evaluations thereof and engineering lessons learned [10, 55, 90, 93]. In contrast, long-term recommender deployments with organic users are rare coming out of academia.

Engineering and maintenance costs, the difficulty of attracting organic usage, and academic incentives are key barriers. A notable exception is MovieLens and other deployed recommenders from the GroupLens academic group [17, 33, 34, 51, 65, 74]; more recent examples include ScholarInbox [22] and POPROX [9, 36]. Most related to our work is ScholarInbox, who build a self-standing platform to recommend research papers, primarily to computer scientists; Flicke et al. [22] describe their system design and survey their users. To our knowledge, PAPER SKYGEST is the first continuously deployed personalized social media feed from academia. To support others to do the same, we contribute an example system design and open source code.

Social Media Feed and Recommender Design. There is substantial academic interest in designing algorithmic feeds [68] – especially those that go beyond engagement optimization in supporting user agency and societal values [14, 40, 48, 56, 57, 59, 81, 82, 88]. However, as Piccardi et al. [68] state, “these questions have remained largely unanswerable to the group of people without personal access to hundreds of millions of users.” More generally, algorithm design for recommenders, including for for academic content, is of substantial interest [26, 35, 46, 49, 66, 67, 69]. However, as Higley et al. [36] discuss, there is a substantial gap between conducting recommender system research and building one. As discussed above, PAPER SKYGEST provides a deployment avenue for such approaches, and more broadly, we aim to help close the gap for academic research and deployments for feed design.

Recent work has also leveraged the opportunities afforded by the decentralized nature of emerging social media [23, 39, 64], primarily by prototyping and evaluating with a small set of (potentially paid) subjects on Bluesky and Mastodon [56, 91]. However, to our knowledge, these systems have not been deployed at scale for organic usage overtime.

Empirical Analyses of Bluesky and Scientific Discovery on Social Media. A number of recent studies have performed observational analyses of the Bluesky ecosystem [5, 19, 42, 72, 80], leveraging its fully open (and monetarily free) data ecosystem. For example,




Figure 2: PAPER SKYGEST usage data over time. (a) Figure 2a shows, for each day since launching, the number of unique accounts that used PAPER SKYGEST and the number of sessions on PAPER SKYGEST (where a session is defined as a request from Bluesky for the first page of posts). The number of unique users has increased over time, while the daily number of sessions has remained stable. Both user and session counts vary throughout the week, with high counts on weekdays and low counts on weekends. (b) Figure 2b shows average usage trajectories for users over time, where users are grouped by their percentile of total PAPER SKYGEST accesses. We consider usage data between May 21, 2025 and September 16, 2025, for users who accessed PAPER SKYGEST for the first time prior to May 21, 2025; we average across users and report 95% confidence intervals.

Quelle and Bovet [71] analyze the large ecosystem of custom feeds. A particular focus has been analyzing how academics have adopted the platform [2, 72, 92]; for example, Quelle et al. [72] identifies over 300,000 academic Bluesky users. More generally, academic behavior on social media – and its implications on science – has been a substantial area of study [1, 6, 13, 15, 24, 28, 45, 47, 61, 83, 85], studying how social media usage and visibility affects scientific outcomes and academic productivity. One notable example is Qiu et al. [70], who experimentally evaluated how social media affects job market outcomes, by assigning EconTwitter influencers to randomly retweet posts about job market papers – women in the treatment group, on average, received 0.9 additional job offers; similarly, Branch et al. [8] experimentally find that papers promoted by leading science communicators were downloaded about $3\times$ more after tweeting and had 81% higher Altmetric scores, but the increase in citations three years later was not statistically significant.

To this literature, we contribute behavioral observations about the effects of social media rankings; for example, a user is $4\times$ more likely to “like” a post that appears first in a their PAPER SKYGEST feed than one that appeared, at highest, fifth. As we discuss, PAPER SKYGEST further provides a deployment setting to experimentally measure such effects, such as through running randomized controlled trials on algorithm design.

3 PAPER SKYGEST AND USAGE

In this section, we give further details on Bluesky’s custom feeds, and describe PAPER SKYGEST. We then provide statistics on our feed’s usage, and show qualitative feedback from Skygest users along with usage statistics. Finally, we illustrate behavioral effects in how users interact with posts and how adopting PAPER SKYGEST corresponded with a shift in behavior.

3.1 Paper Skygest Overview

Custom feeds on Bluesky are recommendation algorithms which are designed and hosted by people external to the Bluesky organization. These feeds are not simply a *feature* Bluesky offers – users are *encouraged* to pin custom feeds as they are a core component of Bluesky’s ecosystem, offering users “algorithmic choice” [50]. When a user discovers a feed she likes – through either her social network or Bluesky’s feed search tool – she can pin it to the feed navigation bar shown in Figure 1a. While using the app, the user can swipe left and right to switch between pinned feeds, and can reorder the feeds in the navigation bar as she chooses.

Under the hood, each time a user navigates to a custom feed such as PAPER SKYGEST in the navigation bar, Bluesky sends the user’s identifier to the custom feed’s *feed generator*, and the feed generator returns a ranked list of posts to show the user. The feed generator’s only responsibility is to determine the list of posts to be shown, and Bluesky places no constraints on this list.

PAPER SKYGEST currently shows a user posts about academic papers from accounts she is following, in reverse-chronological order. PAPER SKYGEST includes a post if its text or metadata include a link to a site hosting preprints or published manuscripts, or contain at least three keywords related to academic paper announcements – we make this logic open-source through GitHub. A more detailed system architecture is provided in Section 4, and we provide a schematic diagram in Figure 1b.

We launched PAPER SKYGEST publicly on March 10, 2025. Between then and September 17, 2025, an average of 1,025 unique users visited PAPER SKYGEST each day, and an average of 1,610 users visited PAPER SKYGEST each week, with growing numbers of active users over time; in total, PAPER SKYGEST was used 1,073,812 times in the time period. To launch PAPER SKYGEST, we simply posted an announcement on Bluesky and encouraged users in our network to spread the word and repost our announcement. Our user base then grew naturally over time as users reposted or recommended

PAPER SKYGEST to other users. Notably, we did not take out advertisements for PAPER SKYGEST or recruit paid participants – our users are organic.

Consent and Research Ethics Approval. This research was approved by our University IRB under Protocol #0149634. For a user’s first five visits, PAPER SKYGEST shows a post thread explaining that we are conducting research, including information about data collection and consent.⁴ Users may request their data to be removed and excluded at any time while still maintaining feed access. In this work, we include a user in our analysis if she completes these five visits and does not opt out of data collection. We exclude our research team from analysis.

In all analyses and usage statistics, we restrict to data prior to September 17, 2025. We exclude users with outlying access patterns (over 20,000 total accesses each), yielding 3,241 PAPER SKYGEST users.

3.2 Quantitative Usage Metrics

PAPER SKYGEST has substantial organic usage. In Figure 2a, we plot usage metrics for PAPER SKYGEST since its launch in March 2025; our user base has increased over time from an average of 885 users per day in April, to an average of 1,144 users per day during August. PAPER SKYGEST has maintained a high level of traffic; on average PAPER SKYGEST is checked 5,622 times across users per day. (If we include users we excluded with over 20,000 daily accesses, PAPER SKYGEST is accessed over 8,000 times per day, for over 1.6 million overall accesses during the studied time period).

PAPER SKYGEST users return over time. PAPER SKYGEST also has many users who have not “churned” off the feed. Each week, an average of 810 users use PAPER SKYGEST for at least five days – a large cohort are dedicated users of the feed. In Figure 2b, we show user trajectories in the 16 weeks after the first time they accessed PAPER SKYGEST, with different cohorts defined by how much they have used PAPER SKYGEST overall. As illustrated in the bottom Blue line, about half of PAPER SKYGEST users (who make it past the 5 consent posts) approximately do not return to the feed after the first week. However, for the remainder groups, average usage 16 weeks after the first access is only slightly lower than average usage in the second week; thus, while the frequency of access differs substantially across users, individual users tend to return to PAPER SKYGEST over time.

3.3 Qualitative Feedback

Since the launch of our feed, we have received substantial interest and positive feedback; almost all user growth stems from word of mouth by PAPER SKYGEST users – we do not conduct online advertising, and to our knowledge, feed discovery avenues on Bluesky are not a big channel for adoption. Here, we analyze qualitative feedback, by users posting about PAPER SKYGEST on Bluesky. The below quotes are organically generated public posts on Bluesky by users independent of the research team, who posted without being asked. We received permission from each quoted user to include their post; we include their handle and link to post if given

⁴This consent information was added several months after initial deployment. When it was added, all current users were shown the consent thread for their next five visits.

explicit permission to do so, and otherwise default to quoting the post anonymously.

Many users emphasize the usefulness of PAPER SKYGEST in highlighting posts about research papers that might otherwise be missed or buried under other content:

- *This might be the most useful thing I have come across in social media - a personalized feed of academic papers filtered by your follower network! Highly recommend. #academicsky - @suryaganguli.bsky.social*
- *This feed really does make bsky useful again for following research. (My home had descended into full politics for the last few weeks.) - @nsaphra.bsky.social*
- *This is the first feed that has felt like proper academic twitter (love it or hate it) - @eugenevinitzsky.bsky.social*
- *This has potential for academics who miss old Econtwitter. - anonymous*
- *This is awesome ! 🍌 One scroll of the feed and I am seeing several papers and posts that I missed via my Discovery feed. - anonymous*

Several users noted in particular that PAPER SKYGEST illustrates the power of externally developed feeds, as benefiting both users and researchers. For example one user, in response to our posts disclosing that we are conducting research, said

I have been thinking for a while about how AT Proto and Bluesky can benefit researchers[.] One way is to help surface relevant papers, which is exactly what Skygest is doing[.] But AT Proto/Bluesky can in itself be used to conduct research, which is also what Skygest is doing! - @o.simardcasanova.net

Another user said

This is ✨excellent✨ as it picks up all the papers that folks mention to each other in replies, as well as those in top level posts. Brilliant way to make use of the power of feeds. - anonymous

3.4 User Behavior

We now leverage data on PAPER SKYGEST accesses and our corpus of posts about academic papers to study how scientific content is shown to and interacted with by PAPER SKYGEST users. These are examples of the analyses that PAPER SKYGEST facilitates – and custom feeds on Bluesky more generally – even before running randomized experiments. We refer to the corpus of posts that pass our filter as “paper posts.”

Engagement as a function of feed ranking. An oft-observed phenomenon in recommender systems and feed design is that item ordering matters – users are more likely to interact with items shown first. For example, Haque and Ginsparg [31, 32] find in 2009 in several physics categories on ArXiv that papers at the top of the daily digest email have over 50% more citations and visibility, with similar positive effects of being last (compared to being in the middle). Similarly, Feenberg et al. [20] found that being listed first in the NBER⁵ email of working papers leads to a 33% increase in views and 27% increase in citations, even though papers are ordered via a process that the email generators view as random.

⁵A prestigious economics organization that sends out working papers by its members




Figure 3: The rate at which users like or repost a post, as a function of how high it appeared in their feed when they switched to it (since users may access the feed multiple times with overlap in posts, we use the highest position that the post appeared). There are strong feed positioning effects. We consider only posts, likes, reposts, and accesses up to September 16, 2025, and filter to user likes and reposts that occurred within 30 seconds of the user accessing PAPER SKYGEST. Bluesky requests pages of size 10 and 30; we use access data for cases where Bluesky requests 30 posts. This results in 5,929 likes and 1,225 reposts. We take the average across post-user pairs for each rank, and report 95% confidence intervals.

To our knowledge, social media companies have not generally released individual-post-and-user level view and interaction data, and so it is unknown how strong these effects are for scientific content on social media. Figure 3 illustrates how users’ rate of interacting with a post changes with the maximum rank the post had when shown to the user in PAPER SKYGEST. We observe a sharp drop-off in likes as the maximum rank increases; while there are far fewer reposts, we observe a similar decrease here. In particular, users are much more likely to interact with a post first in their feed – for example, 4× more likely to “like” a post that appeared first in a their PAPER SKYGEST feed than one that appeared, at highest, fifth. The relative differences are even stronger for reposts. These relationships are much higher than those found in the ArXiv and NBER settings, underscoring the importance of social media algorithm design in mediating scientific consumption. We note that this interaction difference is not directly due to our algorithm prioritizing content that we believe the user will like more, since PAPER SKYGEST currently orders posts in reverse chronological order.⁶

User engagement with scientific posts after adopting PAPER SKYGEST. Next we ask: “does using PAPER SKYGEST correspond to a change in behavior on Bluesky?” We might expect that use of PAPER SKYGEST leads to an increase in engagement with academic content, compared to before the user started using PAPER SKYGEST– since PAPER SKYGEST makes discovery of such posts easier.

⁶There may be other differences explaining the effect, such as that feed ordering and feed usage timing correlates jointly with other reasons a user interacts with posts; for example, that they open our feed immediately after a friend posts a paper thread, to interact with that post. Such confounding can only be mitigated through a randomized controlled trial on feed order; we discuss such experiments in Section 5.

To evaluate this, we compare how each user interacted with posts in the windows of time one week before to one week after the user’s first time accessing PAPER SKYGEST. We query the Bluesky API for all “like” events for PAPER SKYGEST users across all Bluesky posts; we consider users who are already actively engaged with academic content on Bluesky before this window (to mitigate capturing naive effects such as a user first using PAPER SKYGEST immediately after they started using Bluesky). In particular, we consider users who had at least one like of a paper post and one like of a non-paper post in the week before the window (i.e., two weeks to 1 week before they first used PAPER SKYGEST) For technical data collection reasons and to remove potential bots, we further remove users who had more than 10,000 total likes on Bluesky. After these filters, the following analysis corresponds to 796 (25%) PAPER SKYGEST users. (We find directionally similar findings if we make the minimum activity threshold larger, though this removes more users).

For each user, we count the number of likes of paper posts in the week before and the week after first accessing PAPER SKYGEST, $\ell_{\text{paper}}^{\text{before}}(u)$ and $\ell_{\text{paper}}^{\text{after}}(u)$ respectively. We first consider the difference in likes of paper posts in this window after accessing PAPER SKYGEST: the average difference $\ell_{\text{paper}}^{\text{after}}(u) - \ell_{\text{paper}}^{\text{before}}(u)$ is 0.56 with standard error 0.12, yielding a 95% confidence interval of (0.33, 0.78). In other words, on average, a user liked 0.5 more posts about research papers in the week after adopting PAPER SKYGEST than in the week before, and the difference is statistically significant – note that this difference occurs even though “like” interactions are sparse, with users on average liking less than 1% of posts about papers that are shown to them.

One limitation of the above analysis is – even though we filter for users already active on Bluesky – that user activity may be changing over time, and adoption of PAPER SKYGEST may still correlate with an intention to be more active on Bluesky. Thus, to better capture the effect of PAPER SKYGEST, we compare this increase in interactions with the corresponding change in interactions with posts that are *not* about research; if user activity levels are driving the above effect, then we might expect that interactions with non-research posts *also* increase a corresponding amount. Thus, we next analyze whether the *fraction* of a user’s engagement that is with posts about research changes in the week after first using PAPER SKYGEST. Formally, we collect each user’s total likes in the window before and after, $\ell_{\text{total}}^{\text{before}}(u)$ and $\ell_{\text{total}}^{\text{after}}(u)$, and compute the difference in the *proportion* of the user’s likes that are of paper posts

$$\frac{\ell_{\text{paper}}^{\text{after}}(u)}{\ell_{\text{total}}^{\text{after}}(u)} - \frac{\ell_{\text{paper}}^{\text{before}}(u)}{\ell_{\text{total}}^{\text{before}}(u)}.$$

The average difference of proportions is 0.035 with standard error of 0.0078, yielding a 95% confidence interval of (0.019, 0.050). In other words, users do indeed shift their interaction toward posts about research after adopting PAPER SKYGEST, and the effect is small but statistically significant (on average, the percentage of their interactions that are on posts about papers increases by 3.5 percentage points).

Content Distributions. Finally, we compare the academic content shown to users and interacted with on PAPER SKYGEST, compared to the universe of academic content on Bluesky. We are interested

in two questions: (1) do PAPER SKYGEST users come from across scientific disciplines (as opposed to, for example, just computer science); and (2) are there differences across (sub-)disciplines in how users interact with shown content?

First, we find that PAPER SKYGEST users are shown a variety of academic content, reflecting a diverse user base (since content is only shown if a user follows the posting account). We find that 8.9% of posts shown to PAPER SKYGEST users include a link to an article from the general interest journal *Nature* (or a *Nature* subjournal), around 5.8% link to an economics journal or preprint server, and around 14.6% link to a preprint on ArXiv. There is also substantial content from journals belonging to various engineering and humanities disciplines.

Second, we focus on posts that link to ArXiv to study if there are behavioral differences across (sub-)disciplines (since we can retrieve the ArXiv sub-category for each linked paper). For this analysis, we further filter out (do not include) ArXiv posts in our database that are visibly from bot accounts (e.g., accounts that make posts about each new ArXiv paper in a particular category); these bot posts correspond to around 40% of posts with links to ArXiv papers. Of the remaining posts, over 39% are shown to PAPER SKYGEST users at least once. Figure 4 compares the distribution of ArXiv categories across posts in the corpus with the distribution of ArXiv categories across posts shown to PAPER SKYGEST users, and again with the distribution of ArXiv categories across user likes of academic posts. Across the board, there is high representation from computer science and specifically artificial intelligence (AI) and machine learning (ML). Furthermore, posts with papers in AI/ML receive a disproportionate number of likes, compared to either the number of posts on Bluesky in the category, and the number of posts shown to PAPER SKYGEST users. Similarly, posts in the CS.CY category (“Computers and Society”) are both disproportionately shown to PAPER SKYGEST users and disproportionately engaged with upon being shown. On the other hand, posts about quantum physics (quant-ph and gr-gc) are interacted with at lower rates.

4 PAPER SKYGEST ARCHITECTURE

In this section, we describe the backend infrastructure of PAPER SKYGEST, to support others in developing their own feeds. Our code is available at <https://github.com/Skygest/PaperSkygest>. We illustrate our architecture in Figure 1b – primarily built on AWS services. As one measure of implementation challenges others might expect: an initial prototype of the feed, based on a prior template [79],⁷ was built in less than two days in December 2024; however, it had multiple-second latency, and only worked for a few users at a time. Our current architecture, as provided, has around 50 millisecond latency, autoscales with users, and supports multiple algorithms – and was developed over several months and had several test users, before its public launch in March 2025.

Bluesky Custom Feed Generator Overview. A simplified overview of how Bluesky integrates with custom feed generators is as follows; for further technical details, see Bluesky’s documentation [7]. When

a user selects a feed to view (e.g., PAPER SKYGEST), Bluesky⁸ looks up the feed generator (e.g., our system) associated with that feed, and sends a request to that feed generator along with an authentication token which can be used to identify the user. The feed generator determines which posts to show that user, and returns an ordered list of post identifiers – any post on the ATProtocol can be returned, whether or not the user follows the account associated with the post. Bluesky then shows the user (most of) the posts specified in the returned list; it handles all visualization aspects such as hydration of interaction data; it further handles threading, moderation, and post deletions, and so the list of posts presented may not exactly be the list of posts returned by the generator. These steps are hidden from the user – from her perspective, she simply clicks on the feed, and sees a ranked list of posts, all within the Bluesky app.

There are three main components to building a custom feed: (1) handling incoming requests from Bluesky by returning a set of posts, (2) maintaining a corpus of candidate posts to recommend, and (3) implementing recommendation logic determining which posts to show for each user. We give details on how we addressed each of these components for PAPER SKYGEST below.

4.1 Feed Generator Endpoint

A feed generator is a server that hosts a feed; the main requirement of a feed generator is to implement the endpoint `/xrpc/app.bsky.feed.getFeedSkeleton`, which returns a list of posts to show. For PAPER SKYGEST, determining which posts to show has high latency – 75% of our users’ recommendations take over six seconds to generate – since we need to get a user’s follows from the Bluesky API. This is a prohibitively long time for users to wait, so instead we generate recommendations offline and store them in a database, which we then query in real time. Beyond the efficiency gains, decoupling the generation and serving of recommendations also improves robustness – if an error occurs during recommendation generation, the only consequence on the user’s experience is that the user will see stale recommendations.

Given this, PAPER SKYGEST’s implementation of the `getFeedSkeleton` endpoint is as follows: First, we resolve the request’s authentication token to obtain the user’s identifier. If it is the user’s first time, we trigger an event to generate the user’s recommendations offline, and immediately return a single post prompting the user to refresh after a few seconds. Otherwise, if PAPER SKYGEST has seen the user before, we query an AWS DynamoDB database (“recs”) to get a list of cached recommendations for the user (generated as described below in Section 4.3). Before returning posts, we trigger an event to regenerate the user’s recommendations, so that if the cached recommendations are stale, the user can refresh after a few seconds and see new recommendations. We also trigger an `postprocessing` event to asynchronously save the user to a users database, and log the user’s access, including the posts shown. For the first five visits by a user, we further show a consent thread, explaining that we are conducting research and how to opt out. This approach balances low real-time latency with delivering fresh content.

Pagination. The `getFeedSkeleton` request comes with two parameters: `limit` (the number of posts to return in one page of

⁷Our initial code was adapted from an open-source Python template from MarshalX [79], which provides a self-standing single server feed implementation, with a Peewee SQLite database. We decomposed the system into several interdependent modules in AWS EC2 and Lambda and extended it, for auto-scaling, efficiency, and robustness.

⁸In general, custom feeds can be queried by other clients on the ATProtocol; we refer to Bluesky here for simplicity.




Figure 4: The distribution of ArXiv categories among (a) posts on Bluesky containing ArXiv links, (b) all posts containing ArXiv links shown to PAPER SKYGEST users, and (c) likes of posts containing ArXiv links across Bluesky. We show categories that are in the top 10 most common categories for at least one of these distributions. We limit our consideration to posts from non-bot accounts. We identify bots as accounts which mention arxiv or bot in their handle, or have very high posting activity on Bluesky. We use public data from ArXiv [3] to find a post’s categories given its ArXiv identifier. There is heterogeneity in the rate at which posts are interacted with, across categories.

content) and, if the user has already consumed a page of content, a cursor (the current page number). We return the posts from index cursor to cursor + limit − 1 from the list of posts queried from the recs database, along with cursor + limit as the new cursor.

Cloud Hosting Details. We host this logic on AWS Lambda, a serverless compute service which scales allocated compute resources with usage. This allows us to adapt and scale to different system loads across time in a robust and cost-effective manner. We set up the endpoint on AWS API Gateway, and attach it to the Lambda function. The postprocessing occurs in a separate Lambda function. The recommendation generation and postprocessing events are triggered using AWS SQS. If the Lambda function is not called for a few minutes, it will shut down and restart the next time it is invoked; since restarting introduces a delay in serving recommendations, we set up an AWS EventBridge schedule to invoke the function every four minutes to keep the Lambda function active.

In addition to the getFeedSkeleton endpoint, we also publish a feed record with a unique identifier, and implement two other required endpoints (/xrpc/app.bsky.feed.describeFeedGenerator and /.well-known/did.json). We again implement these endpoints using AWS Lambda and API Gateway; the code for these two endpoints does not differ significantly from the template [79].

4.2 Firehose

To assemble a corpus of Bluesky posts, we set up an AWS EC2 server to subscribe to the Bluesky firehose. Put simply, the firehose is a stream of public “events” on Bluesky, such as a user creating or liking a post, or following another user [50]. In particular, we listen for post creation events. We examine each new post’s text and metadata (including embedded links), and filter to posts that are about an academic paper – namely, posts that either (a) contain a link to an academic journal or platforms which host preprints and working papers, or (b) have at least three keywords relating to announcing new academic work. We also make this filter open-source.

We store these posts in an AWS DynamoDB database (“posts”), and maintain an index of post authors sorted by creation date.

In addition to post creations, we also listen for several other events on the firehose, including user interactions (for example, a user likes a post). We maintain a local database on EC2 of the unique identifiers of all posts about papers in our dataset, and filter to interactions with paper posts. We save these filtered interactions to additional DynamoDB databases. We also listen for post deletions to mark these posts as deleted in our posts database.

For efficiency, after ingesting relevant events, we add the event to a queue. We set up parallel worker processes to take events from the queue to filter and write to the database.

4.3 Recommendation Generation

The recommendation generation process is what is used to calculate recommendations for each user, given the database of paper posts and a set of recommendation algorithms. It is decomposed into two components: the dispatch module, which assigns PAPER SKYGEST users to batches, and the generation module, which generates recommendations for each batch in parallel.

Recommendation Dispatch Module. The recommendation dispatch module queries the users database for a list of PAPER SKYGEST users. The module randomly assigns these users into batches of 20, and triggers an asynchronous invocation of the generation module for each batch. This module is scheduled to run every 20 minutes.

Recommendation Generation Module. For each user in a batch, this module generates recommendations and saves them to the recs database, the database queried by the feed generator queries when a user requests the feed.

For each user, the generation module runs one or more recommendation algorithms. At the moment, the vast majority of users are shown recommendations from the “reverse chronological” algorithm, which simply shows posts from accounts the user follows

in reverse chronological order. Given a user’s identifier, we request the accounts that user is following from Bluesky’s public API. Then, for each of these accounts, we query the author index of posts to find the ten most recent posts from that account. We then sort all of these posts chronologically, and store the first 150 in recs. We note that we query the users’ followers at every recommendation generation; the alternative would be to maintain the follower graph (which can be done by listening to the appropriate events), which would be comparatively expensive computationally and less robust.

At generation time, we also currently generate alternative recommendations that include reposts and quote posts from a user’s follows; these feeds are in deployment for several beta testers. To deploy such alternative recommendations, our feed generator endpoint simply needs to request a different algorithm from the recs database; indeed, the endpoint could request a different algorithm for each user. We also save the first 30 posts from each ranking to a `counterfactual_recs` database to facilitate future analyses, so we can see what a user would have viewed if she had checked the feed at that time and if we had used alternative algorithms, unless the user has requested removal from research consideration; this allows us to conduct the type of counterfactual comparison research suggested by González-Bailón et al. [25].

Finally, since the feed generator endpoint only interfaces with the recs database, we can also split the recommendation generation into independent submodules for different algorithms. We are currently developing a prototype of an algorithmic feed; when it is deployed we will use host the recommendation generation logic in a separate module from the following feed.

Cloud Hosting Details. The generation and dispatch modules are each hosted on separate AWS Lambda functions. The dispatch module is triggered by an AWS EventBridge schedule, and sends generation messages to the generation module using Amazon SQS. Both recs and `counterfactual_recs` are AWS DynamoDB databases.

Cold Start and Onboarding. If the user has checked the feed fewer than ten times at the time of recommendation generation, we prepend a post with details about PAPER SKYGEST, (along with potentially the thread with consent information).

When a user’s Skygest feed contains fewer than 10 posts (because the accounts they follow have not posted enough research content), we append a post encouraging the user to follow more accounts, and then a “default” feed, generated as follows. During the dispatch stage, we also randomly select one of the batches for regenerating the default feed. When we process that batch of recommendations in the recommendations module, we save the most recent 150 posts across all of the posts recommended to users in that batch. We also show this default feed to logged-out users, since we do not have followers for such users.

5 APPLICATIONS OF ATPROTOCOL AND PAPER SKYGEST

Custom feeds on Bluesky are a new avenue for academics to build, deploy, and evaluate social media algorithms with organic users, without requiring a partnership with the platform itself. Our experience with PAPER SKYGEST demonstrates that academics can

use custom feeds to host recommenders that garner sustained usage from organic users. We believe that custom feeds, and PAPER SKYGEST in particular, enable study of new questions of interest to the community. We further hope that our architecture description, along with provided code, supports others in this analysis.

Causally Studying the Effects of Algorithm Design. While the results we present above are observations of user behavior, the infrastructure we built for PAPER SKYGEST is designed to facilitate experimentation (randomized controlled trials). We support using different algorithms for different users, including randomizing across users. We also save counterfactual recommendations for each user across multiple different recommendation algorithms, and collect public interaction data like user likes and reposts over time. We can also insert posts soliciting survey responses in the same way we inserted posts with consent information at the beginning of the feed. With these ingredients, we can compare user engagement and satisfaction across feed designs. (We note that comparing feed designs at the post level will be challenging, due to ranking order effects, cf. Schein et al. [76]).

Designing for User Agency. While custom feeds provide users algorithmic choice through selecting from a pre-defined set of algorithms, recent work has proposed and explored platform designs where users have fine-grained control over their recommendations [60], even *within* a feed [40, 41]. For example, Carroll et al. [11] introduced a system to steer recommendations through natural language, Instagram launched a page for users to select topics to include in their feeds [38], and Malki et al. [56] built an interface with a wide range of controls. Bluesky provides an exciting opportunity to deploy controllable feeds in particular because Bluesky users, already familiar with algorithmic choice, may be more interested in organically adopting controllable feeds. One avenue is building an external interface where users can configure their preferences for a custom feed like PAPER SKYGEST (e.g. tuning parameters such as the weight on reposts or content diversity), which are then used to populate the custom feed for that user. Malki et al. [56] make progress by piloting a user control interface for Bluesky posts, but do not yet integrate it with a deployed custom feed.

Feed Opportunities Beyond PAPER SKYGEST. Beyond PAPER SKYGEST, a wide variety of interesting open questions exist. For example, feeds with polarizing content could be used to explore novel *bridging* algorithms [89]. Feeds focused on entertainment could be used to study overconsumption and related interventions – e.g., one could explore the effect of a delay or a wellness reminder before a user is shown the feed, as in Grüning et al. [27]. More broadly – since the custom feed design does not restrict which posts can be shown, and all Bluesky posts can be ingested via the firehose – the space of algorithm design is large, in contrast to a commonly faced challenge in browser-extension based re-ranking [68].

Limitations and User Privacy. One limitation of custom feed experiments is that these experiments may have low external validity due to idiosyncracies in Bluesky’s or the feed’s user population. Furthermore, while the ATPProtocol is open by design, the trade-off with user privacy is an ongoing conversation. For example, Bluesky does not make fine-grained user analytics – such as click and dwell time data – available to custom feed owners: historically,

the only non-public data available to feed designers is the request from Bluesky to return a list of posts (indicating a user access), and which posts the feed generator returned. Thus, researchers are largely restricted to using the coarser, public interaction data to train and evaluate recommenders. One exception is that Bluesky recently started providing feed owners the data from when a user clicks the “see less/more like this” button on their custom feed. Finally, while decentralized social media platforms are an exciting opportunity for academics to build natively on-platform, they still need to acquire organic users by building products users want.

6 CONCLUSION

In this work, we present PAPER SKYGEST, a custom feed on Bluesky showing personalized posts about papers for academic users. We show quantitative evidence of sustained usage and positive qualitative feedback for PAPER SKYGEST, and use data from PAPER SKYGEST to explore academic Bluesky usage. We provide a detailed overview of our infrastructure to allow others to learn from our deployment experience, and share our full code. Finally, we highlight the potential of custom feeds such as PAPER SKYGEST for facilitating academic research on recommendation algorithms using organic users and outside of industry partnerships.

REFERENCES

- [1] Eleonora Alabrese, Francesco Capozza, and Prashant Garg. 2024. Politicized scientists: Credibility cost of political expression on Twitter. (2024). 3
- [2] Wenceslao Arroyo-Machado, Nicolas Robinson-Garcia, and Daniel Torres-Salinas. 2025. Are there stars in Bluesky? A comparative exploratory analysis of altmetric mentions between X and Bluesky. *Journal of Informetrics* 19, 3 (2025), 101700. <https://doi.org/10.1016/j.joi.2025.101700> 3
- [3] arXiv.org submitters. 2024. arXiv Dataset. <https://doi.org/10.34740/KAGGLE/DSV/7548853> 7
- [4] Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. 1
- [5] Leonhard Balduf, Saidu Sokoto, Onur Ascigil, Gareth Tyson, Björn Scheuermann, Maciej Korczyński, Ignacio Castro, and Michaundefined Król. 2024. Looking AT the Blue Skies of Bluesky. In *Proceedings of the 2024 ACM on Internet Measurement Conference* (Madrid, Spain) (IMC '24). Association for Computing Machinery, New York, NY, USA, 76–91. <https://doi.org/10.1145/3646547.3688407> 2
- [6] Holly M Bik and Miriam C Goldstein. 2013. An introduction to social media for scientists. *PLoS biology* 11, 4 (2013), e1001535. 3
- [7] Bluesky, PBC. [n.d.]. Custom Feeds. <https://docs.bsky.app/docs/starter-templates/custom-feeds> 6
- [8] Trevor A Branch, Isabelle M Côté, Solomon R David, Joshua A Drew, Michelle LaRue, Melissa C Márquez, E Chris M Parsons, D Rabaiotti, David Shiffman, David A Steen, et al. 2024. Controlled experiment finds no detectable citation bump from Twitter promotion. *Plos one* 19, 3 (2024), e0292201. 3
- [9] Robin Burke and Michael Ekstrand. 2025. Conducting Recommender Systems User Studies Using POPROX. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*. 1–2. 2
- [10] Qingpeng Cai, Zhenghai Xue, Chi Zhang, Wanqi Xue, Shuchang Liu, Ruohan Zhan, Xueliang Wang, Tianyou Zuo, Wentao Xie, Dong Zheng, et al. 2023. Two-stage constrained actor-critic for short video recommendation. In *Proceedings of the ACM web conference 2023*. 865–875. 2
- [11] Micah Carroll, Adeline Foote, Marcus Williams, Anca Dragan, W. Bradley Knox, and Smitha Milli. 2025. CTRL-Rec: Controlling Recommender Systems With Natural Language. In *ICLR 2025 Workshop on Bidirectional Human-AI Alignment*. <https://openreview.net/forum?id=IX1gdoAoV> 8
- [12] Sarah H. Cen, Andrew Ilyas, Jennifer Allen, Hannah Li, and Aleksander Madry. 2024. Measuring Strategization in Recommendation: Users Adapt Their Behavior to Shape Future Content. In *Proceedings of the 25th ACM Conference on Economics and Computation* (New Haven, CT, USA) (EC '24). Association for Computing Machinery, New York, NY, USA, 203–204. <https://doi.org/10.1145/3670865.3673634> 1
- [13] Kimberley Collins, David Shiffman, and Jenny Rock. 2016. How are scientists using social media in the workplace? *PloS one* 11, 10 (2016), e0162680. 3
- [14] Tom Cunningham, Sana Pandey, Leif Sigerson, Jonathan Stray, Jeff Allen, Bonnie Barrilleaux, Ravi Iyer, Mohit Kothari, Behnam Rezaei, Sanjay Kairam, et al. 2025. Ranking by engagement and non-engagement signals: Learnings from industry. *Annals of the New York Academy of Sciences* (2025). 2
- [15] Emily S Darling, David Shiffman, Isabelle M Côté, and Joshua A Drew. 2013. The role of Twitter in the life cycle of a scientific publication. *arXiv preprint arXiv:1305.0435* (2013). 3
- [16] Dominic DiFranzo and Natalie N. Bazarova. 2018. The Truman Platform: Social media simulation for experimental research. (2018). 1
- [17] Michael D. Ekstrand, Daniel Kluver, F. Maxwell Harper, and Joseph A. Konstan. 2015. Letting Users Choose Recommender Algorithms: An Experimental Study. In *Proceedings of the 9th ACM Conference on Recommender Systems* (Vienna, Austria) (RecSys '15). Association for Computing Machinery, New York, NY, USA, 11–18. <https://doi.org/10.1145/2792838.2800195> 2
- [18] Ziv Epstein and Hause Lin. 2022. Yourfeed: Towards open science and interoperable systems for social media. *arXiv preprint arXiv:2207.07478* (2022). 1
- [19] Andrea Failla and Giulio Rossetti. 2024. “I’m in the Bluesky Tonight”: Insights from a year worth of social data. *PLOS ONE* 19, 11 (11 2024), 1–17. <https://doi.org/10.1371/journal.pone.0310330> 2
- [20] Daniel Feenberg, Ina Ganguli, Patrick Gaulé, and Jonathan Gruber. 2017. It’s Good to Be First: Order Bias in Reading and Citing NBER Working Papers. *The Review of Economics and Statistics* 99, 1 (03 2017), 32–39. 4
- [21] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, S1 (2016), 298–320. 1
- [22] Markus Flicke, Glenn Angrabeit, Madhav Iyengar, Vitalii Protchenko, Illia Shakun, Jovan Cicvaric, Bora Kargi, Haoyu He, Lukas Schuler, Lewin Scholz, et al. 2025. Scholar Inbox: Personalized Paper Recommendations for Scientists. *arXiv preprint arXiv:2504.08385* (2025). 2
- [23] Francis Fukuyama, Barak Richman, Ashish Goel, Roberta R Katz, A Douglas Melamed, and Marietje Schaake. 2020. Middleware for dominant digital platforms: A technological solution to a threat to democracy. *CyberPolicy Center, Freeman Spogli Institute* (2020). 1, 2
- [24] Prashant Garg and Thimo Fetzter. 2025. Political expression of academics on social media. *Nature Human Behaviour* (2025). 3
- [25] Sandra González-Bailón, David Lazer, Pablo Barberá, Meiqing Zhang, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Deen Freelon, Matthew Gentzkow, Andrew M Guess, et al. 2023. Asymmetric ideological segregation in exposure to political news on Facebook. *Science* 381, 6656 (2023), 392–398. 1, 8
- [26] Sophie Greenwood, Sudalakshmei Chiniah, and Nikhil Garg. 2024. User-item fairness tradeoffs in recommendations. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 114236–114288. https://proceedings.neurips.cc/paper_files/paper/2024/file/cf836ef32fd53493e02d26670f04d46-Paper-Conference.pdf 2
- [27] David J. Grüning, Frederik Riedel, and Philipp Lorenz-Spreen. 2023. Directing smartphone use through the self-nudge app one sec. *Proceedings of the National Academy of Sciences* 120, 8 (2023), e2213114120. <https://doi.org/10.1073/pnas.2213114120> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2213114120> 8
- [28] Lars Guenther, Claudia Wilhelm, Corinna Oschatz, and Janise Brück. 2023. Science communication on Twitter: Measuring indicators of engagement and their links to user interaction in communication scholars’ Tweet content. *Public Understanding of Science* 32, 7 (2023), 860–869. 3
- [29] Andrew M. Guess, Neil Malhotra, Jennifer Pan, Pablo Barberá, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Edward Kennedy, Young Mie Kim, David Lazer, Devra Moehler, Brendan Nyhan, Carlos Velasco Rivera, Jaime Settle, Daniel Robert Thomas, Emily Thorson, Rebekah Tromble, Arjun Wilkins, Magdalena Wojcieszak, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker. 2023. How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science* 381, 6656 (2023), 398–404. <https://doi.org/10.1126/science.abp9364> arXiv:<https://www.science.org/doi/pdf/10.1126/science.abp9364> 1
- [30] Andrew M. Guess, Neil Malhotra, Jennifer Pan, Pablo Barberá, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Edward Kennedy, Young Mie Kim, David Lazer, Devra Moehler, Brendan Nyhan, Carlos Velasco Rivera, Jaime Settle, Daniel Robert Thomas, Emily Thorson, Rebekah Tromble, Arjun Wilkins, Magdalena Wojcieszak, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker. 2023. Reshares on social media amplify political news but do not detectably affect beliefs or opinions. *Science* 381, 6656 (2023), 404–408. <https://doi.org/10.1126/science.add8424> arXiv:<https://www.science.org/doi/pdf/10.1126/science.add8424> 1
- [31] Asif-ul Haque and Paul Ginsparg. 2009. Positional effects on citation and readership in arXiv. *Journal of the American Society for Information Science and Technology* 60, 11 (2009), 2203–2218. <https://doi.org/10.1002/asi.21166> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.21166> 4
- [32] Asif-ul Haque and Paul Ginsparg. 2010. Last but not least: Additional positional effects on citation and readership in arXiv. *Journal of the American Society for*

- Information Science and Technology 61, 12 (2010), 2381–2388. 4
- [33] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. <https://doi.org/10.1145/2827872.2>
 - [34] F. Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen. 2015. Putting Users in Control of their Recommendations. In *Proceedings of the 9th ACM Conference on Recommender Systems* (Vienna, Austria) (*RecSys '15*). Association for Computing Machinery, New York, NY, USA, 3–10. <https://doi.org/10.1145/2792838.2800179.2>
 - [35] Andreas Haupt, Dylan Hadfield-Menell, and Chara Podimata. 2023. *Recommending to Strategic Users*. Papers 2302.06559. arXiv.org. <https://doi.org/None.2>
 - [36] Karl Higley, Robin Burke, Michael D Ekstrand, and Bart P Knijnenburg. 2025. What News Recommendation Research Did (But Mostly Didn't) Teach Us About Building A News Recommender. *arXiv preprint arXiv:2509.12361* (2025). 2
 - [37] Ferenc Huszár, Sofia Ira Ktena, Conor O'Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2022. Algorithmic amplification of politics on Twitter. *Proceedings of the national academy of sciences* 119, 1 (2022), e2025334119. 1
 - [38] Andrew Hutchinson. 2025. Instagram Tests Reels Algorithm Control Options. <https://www.socialmediatoday.com/news/instagram-rolls-out-reels-topic-algorithm-control-option/802426/> 8
 - [39] Sohyeon Hwang, Sophie Rollins, Thatiany Andrade Nunes, Yuhan Liu, Richmond Wong, Aaron Shaw, and Andrés Monroy-Hernández. 2025. Governing Together: Toward Infrastructure for Community-Run Social Media. *arXiv preprint arXiv:2509.19653* (2025). 2
 - [40] Farnaz Jahanbakhsh, Dora Zhao, Tiziano Piccardi, Zachary Robertson, Ziv Epstein, Sanmi Koyejo, and Michael S Bernstein. 2025. Value Alignment of Social Media Ranking Algorithms. *arXiv preprint arXiv:2509.14434* (2025). 1, 2, 8
 - [41] Dietmar Jannach, Sidra Naveed, and Michael Jugovac. 2017. User Control in Recommender Systems: Overview and Interaction Challenges. *Lecture Notes in Business Information Processing* 278, 21–33. https://doi.org/10.1007/978-3-319-53676-7_2.8
 - [42] Ujun Jeong, Bohan Jiang, Zhen Tan, H. Russell Bernard, and Huan Liu. 2024. Descriptor: A Temporal Multi-network Dataset of Social Interactions in Bluesky Social (BlueTempNet). *IEEE Data Descriptions* 1 (2024), 71–79. <https://doi.org/10.1109/IEEEDATA.2024.3474640.2>
 - [43] Chenyan Jia, Michelle S Lam, Minh Chau Mai, Jeffrey T Hancock, and Michael S Bernstein. 2024. Embedding democratic values into social media AIs via societal objective functions. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–36. 1
 - [44] Sebastian Jilke, Jiahuan Lu, Chengxin Xu, and Shugo Shinohara. 2019. Using large-scale social media experiments in public administration: Assessing charitable consequences of government funding of nonprofits. *Journal of Public Administration Research and Theory* 29, 4 (2019), 627–639. 1
 - [45] Jakob Jünger and Birte Fährnich. 2020. Does really no one care? Analyzing the public engagement of communication scientists on Twitter. *New Media & Society* 22, 3 (2020), 387–408. 3
 - [46] Hyeonsu B Kang, Rafal Kocielnik, Andrew Head, Jiangjiang Yang, Matt Latzke, Aniket Kittur, Daniel S Weld, Doug Downey, and Jonathan Bragg. 2022. From Who You Know to What You Read: Augmenting Scientific Recommendations with Implicit Social Networks. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 302, 23 pages. <https://doi.org/10.1145/3491102.3517470.2>
 - [47] Qing Ke, Yong-Yeol Ahn, and Cassidy R Sugimoto. 2017. A systematic identification and analysis of scientists on Twitter. *PLoS one* 12, 4 (2017), e0175368. 3
 - [48] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2024. The challenge of understanding what users want: Inconsistent preferences and engagement optimization. *Management science* 70, 9 (2024), 6336–6355. 2
 - [49] Jon Kleinberg, Emily Ryu, and Éva Tardos. 2024. Calibrated Recommendations for Users with Decaying Attention. *arXiv:2302.03239 [cs.DS]* <https://arxiv.org/abs/2302.03239.2>
 - [50] Martin Kleppmann, Paul Frazee, Jake Gold, Jay Graber, Daniel Holmgren, Devin Ivy, Jeromy Johnson, Bryan Newbold, and Jaz Volpert. 2024. Bluesky and the AT Protocol: Usable Decentralized Social Media. In *Proceedings of the ACM Conext 2024 Workshop on the Decentralization of the Internet* (Los Angeles, CA, USA) (*DIN '24*). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3694809.3700740.1.3.7>
 - [51] Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. 1997. Grouplens: Applying collaborative filtering to usenet news. *Commun. ACM* 40, 3 (1997), 77–87. 2
 - [52] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24 (2014), 8788–8790. 1
 - [53] David Lazer. 2020. Studying human attention on the Internet. *Proceedings of the National Academy of Sciences* 117, 1 (2020), 21–22. 1
 - [54] Hause Lin, Haritz Garro, Nils Wernerfelt, Jesse Shore, Adam Hughes, Daniel Deisenroth, Nathaniel Barr, Adam Berinsky, Dean Eckles, Gordon Pennycook, et al. 2024. Reducing misinformation sharing at scale using digital accuracy prompt ads. *PsyArXiv* (2024). 1
 - [55] Hongzun Liu, Kang Yin, Tianyu Sun, Rui Huang, Yunsong Li, Xiao Fang, Zhaojie Liu, Weidong Liu, and Guorui Zhou. 2024. De-Anchor: Mitigating Attention Polarization for Lifelong User Behavior Modeling in Click-Through Rate Prediction. In *Companion Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) (*WWW '24*). Association for Computing Machinery, New York, NY, USA, 694–697. <https://doi.org/10.1145/3589335.3651486.2>
 - [56] Omar El Malki, Marianne Aubin Le Quéré, Andrés Monroy-Hernández, and Manoel Horta Ribeiro. 2025. Bonsai: Intentional and Personalized Social Media Feeds. *arXiv preprint arXiv:2509.10776* (2025). 2, 8
 - [57] Smitha Milli, Luca Belli, and Moritz Hardt. 2021. From Optimizing Engagement to Measuring Value. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (*FAccT '21*). Association for Computing Machinery, New York, NY, USA, 714–722. <https://doi.org/10.1145/3442188.3445933.2>
 - [58] Smitha Milli, Micah Carroll, Yike Wang, Sashrika Pandey, Sebastian Zhao, and Anca D Dragan. 2025. Engagement, user satisfaction, and the amplification of divisive content on social media. *PNAS nexus* 4, 3 (2025), pfa062. 1
 - [59] Smitha Milli, Emma Pierson, and Nikhil Garg. 2023. Choosing the right weights: Balancing value, strategy, and noise in recommender systems. *arXiv preprint arXiv:2305.17428* (2023). 2
 - [60] Alex Moehring, Alissa Cooper, Arvind Narayanan, Aviv Ovadya, Elissa Redmiles, Jeff Allen, Jonathan Stray, Julia Kamin, Leif Sigerson, Luke Thorburn, Matt Motyl, Motahare Eslami, Nadine Farid Johnson, Nathaniel Lubin, Ravi Iyer, and Zander Arnao. 2025. Better Feeds: Algorithms That Put People First. *Knight-Georgetown Institute* (2025). <https://kgi.georgetown.edu/research-and-commentary/better-feeds/> 8
 - [61] Philippe Mongeon, Timothy D Bowman, and Rodrigo Costas. 2023. An open data set of scholars on Twitter. *Quantitative Science Studies* 4, 2 (2023), 314–324. 3
 - [62] Mohsen Mosleh, Gordon Pennycook, and David G Rand. 2022. Field experiments on social media. *Current Directions in Psychological Science* 31, 1 (2022), 69–75. 1
 - [63] Brendan Nyhan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, et al. 2023. Like-minded sources on Facebook are prevalent but not polarizing. *Nature* 620, 7972 (2023), 137–144. 1
 - [64] Tolulope Oshinowo, Sohyeon Hwang, Amy X Zhang, and Andrés Monroy-Hernández. 2025. Seeing the Politics of Decentralized Social Media Protocols. *arXiv preprint arXiv:2505.22962* (2025). 2
 - [65] Katherine Panciera, Reid Priedhorsky, Thomas Erickson, and Loren Terveen. 2010. Lurking? Cyclopaths? A quantitative lifecycle analysis of user behavior in a geowiki. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1917–1926. 2
 - [66] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*. 1929–1942. 2
 - [67] Kenny Peng, Manish Raghavan, Emma Pierson, Jon Kleinberg, and Nikhil Garg. 2024. Reconciling the accuracy-diversity trade-off in recommendations. In *Proceedings of the ACM Web Conference 2024*. 1318–1329. 2
 - [68] Tiziano Piccardi, Martin Saveski, Chenyan Jia, Jeffrey Hancock, Jeanne L Tsai, and Michael S Bernstein. 2024. Reranking social media feeds: A practical guide for field experiments. *arXiv preprint arXiv:2406.19571* (2024). 1, 2, 8
 - [69] Tiziano Piccardi, Martin Saveski, Chenyan Jia, Jeffrey T. Hancock, Jeanne L. Tsai, and Michael Bernstein. 2024. Social Media Algorithms Can Shape Affective Polarization via Exposure to Antidemocratic Attitudes and Partisan Animosity. *arXiv:2411.14652 [cs.CY]* <https://arxiv.org/abs/2411.14652.1.2>
 - [70] Jingyi Qiu, Yan Chen, Alain Cohn, and Alvin E Roth. 2024. Social media and job market success: a field experiment on Twitter. *Available at SSRN 4778120* (2024). 3
 - [71] Dorian Quelle and Alexandre Bovet. 2025. Bluesky: Network topology, polarization, and algorithmic curation. *PLoS ONE* (2025). <https://doi.org/10.1371/journal.pone.0318034.3>
 - [72] Dorian Quelle, Frederic Denker, Prashant Garg, and Alexandre Bovet. 2025. Why Academics Are Leaving Twitter for Bluesky. *arXiv preprint arXiv:2505.24801* (2025). 1, 2, 3
 - [73] M.J. Salganik, P.S. Dodds, and D.J. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (2006), 854–6. <https://doi.org/10.1126/science.1121066.1>
 - [74] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web* (Hong Kong, Hong Kong) (*WWW '01*). Association for Computing Machinery, New York, NY, USA, 285–295. <https://doi.org/10.1145/371920.372071.2>
 - [75] Martin Saveski, Nabeel Gillani, Ann Yuan, Prashanth Vijayaraghavan, and Deb Roy. 2022. Perspective-taking to reduce affective polarization on social media. In

- Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 885–895. 1
- [76] Aaron Schein, Keyon Vafa, Dhanya Sridhar, Victor Veitch, Jeffrey Quinn, James Moffet, David M Blei, and Donald P Green. 2021. Assessing the effects of friend-to-friend texting on turnout in the 2018 us midterm elections. In *Proceedings of the Web Conference 2021*. 2025–2036. 8
 - [77] Elisa Shearer. 2021. More than eight-in-ten Americans get news from digital devices. *Pew Research Center* 12 (2021). 1
 - [78] D S Shiffman and J Wester. 2025. Scientists no Longer Find Twitter Professionally Useful, and have Switched to Bluesky. *Integrative and Comparative Biology* (07 2025), icaf127. <https://doi.org/10.1093/icb/ica127> arXiv:<https://academic.oup.com/icb/advance-article-pdf/doi/10.1093/icb/ica127/63723791/ica127.pdf> 1
 - [79] Ilya Siamionau. 2025. bluesky-feed-generator. <https://github.com/MarshallX/bluesky-feed-generator> 6, 7
 - [80] Alyssa Smith, Ilya Amburg, Sagar Kumar, Brooke Foucault Welles, and Nicholas W. Landry. 2025. A Blue Start: A large-scale pairwise and higher-order social network dataset. arXiv:2505.11608 [physics.soc-ph] <https://arxiv.org/abs/2505.11608> 2
 - [81] Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Chloe Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, et al. 2024. Building human values into recommender systems: An interdisciplinary synthesis. *ACM Transactions on Recommender Systems* 2, 3 (2024), 1–57. 2
 - [82] Jonathan Stray, Ivan Vendrov, Jeremy Nixon, Steven Adler, and Dylan Hadfield-Menell. 2021. What are you optimizing for? aligning recommender systems with human values. *arXiv preprint arXiv:2107.10939* (2021). 2
 - [83] Leona Yi-Fan Su, Dietram A Scheufele, Larry Bell, Dominique Brossard, and Michael A Xenos. 2017. Information-sharing and community-building: Exploring the use of Twitter in science public relations. *Science Communication* 39, 5 (2017), 569–597. 3
 - [84] Mike Taylor and Carlos Areia. 2025. Bluesky’s ahead, but is X a dead parrot? <https://www.altmetric.com/blog/blueskys-ahead-but-is-x-a-dead-parrot/> 1
 - [85] Laura Van Eperen and Francesco M Marincola. 2011. How scientists use social media to communicate their research. *Journal of Translational Medicine* 9, 1 (2011), 199. 3
 - [86] Angelina Voggenreiter, Sophie Brandt, Fabian Putterer, Andreas Frings, and Juergen Pfeffer. 2024. The Role of Likes: How Online Feedback Impacts Users’ Mental Health. In *Proceedings of the 16th ACM Web Science Conference*. 302–310. 1
 - [87] Brian E Weeks, Alberto Ardèvol-Abreu, and Homero Gil de Zúñiga. 2017. On-line influence? Social media use, opinion leadership, and political persuasion. *International journal of public opinion research* 29, 2 (2017), 214–239. 1
 - [88] E Glen Weyl, Luke Thorburn, Emillie de Keulenaar, Jacob Mchangama, Divya Siddarth, and Audrey Tang. 2025. Prosocial Media. *arXiv preprint arXiv:2502.10834* (2025). 2
 - [89] Stefan Wojcik, Sophie Hilgard, Nick Judd, Delia Mocanu, Stephen Ragain, M. B. Fallin Hunzaker, Keith Coleman, and Jay Baxter. 2022. Birdwatch: Crowd Wisdom and Bridging Algorithms can Inform Understanding and Reduce the Spread of Misinformation. arXiv:2210.15723 [cs.SI] <https://arxiv.org/abs/2210.15723> 8
 - [90] Xinyi Wu, Donald Loveland, Runjin Chen, Yozen Liu, Xin Chen, Leonardo Neves, Ali Jadbabaie, Mingxuan Ju, Neil Shah, and Tong Zhao. 2025. GraphHash: Graph Clustering Enables Parameter Efficiency in Recommender Systems. In *Proceedings of the ACM on Web Conference 2025*. 357–369. 2
 - [91] Owen Xingjian Zhang, Sohyeon Hwang, Yuhao Liu, Manoel Horta Ribeiro, and Andrés Monroy-Hernández. 2025. Understanding Community-Level Blocklists in Decentralized Social Media. *arXiv preprint arXiv:2506.05522* (2025). 2
 - [92] Er-Te Zheng, Xiaorui Jiang, Zhichao Fang, and Mike Thelwall. 2025. How is science discussed on Bluesky? *arXiv preprint arXiv:2507.18840* (2025). 3
 - [93] Ruitao Zhu, Detao Lv, Yao Yu, Ruihao Zhu, Zhenzhe Zheng, Ke Bu, Quan Lu, and Fan Wu. 2023. LInet: A Location and Intention-Aware Neural Network for Hotel Group Recommendation. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW ’23). Association for Computing Machinery, New York, NY, USA, 779–789. <https://doi.org/10.1145/3543507.3583202> 2