

Saving the Web

Mat Kelly

mkelly@drexel.edu

Assistant Professor, Information Science
Drexel University College of Computing and Informatics

INF0591 Week 10

December 4, 2025

Supported by:

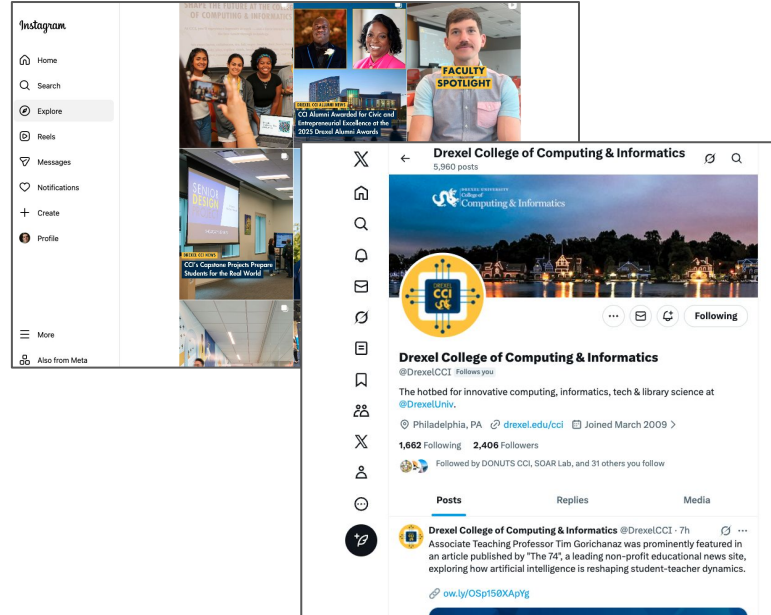
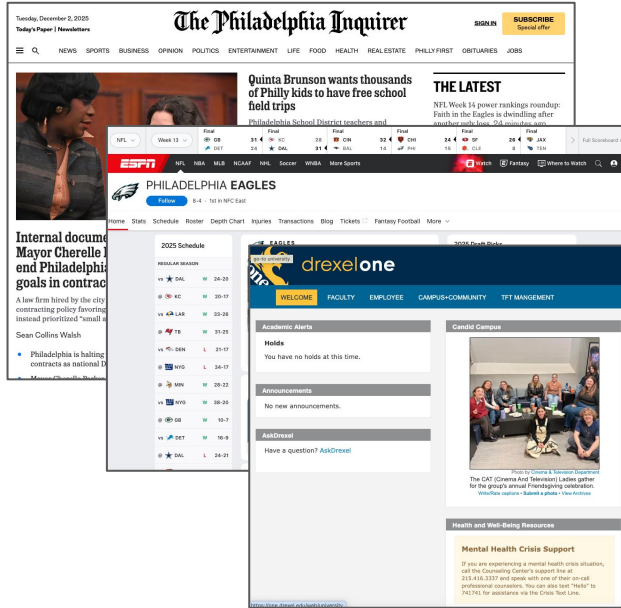
INSTITUTE of
Museum and Library
SERVICES



Slides: <https://bit.ly/info591>

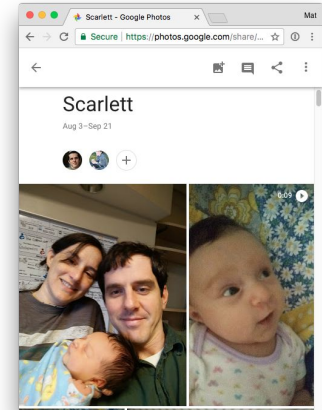
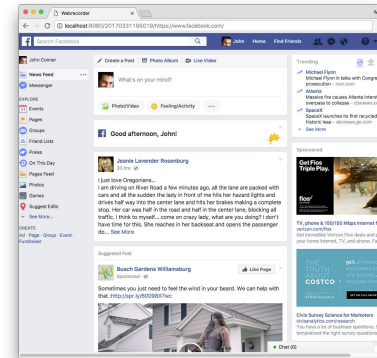
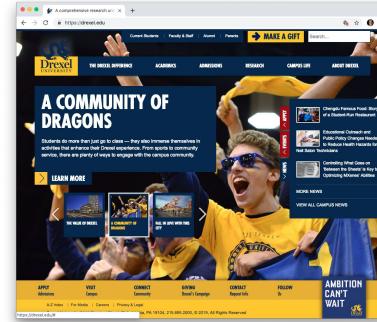
<https://matkelly.com>

The Web



The Web's Importance

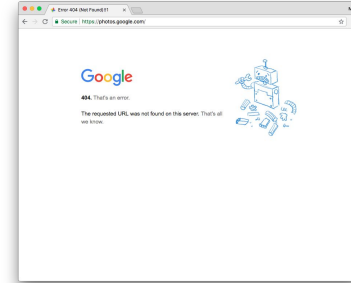
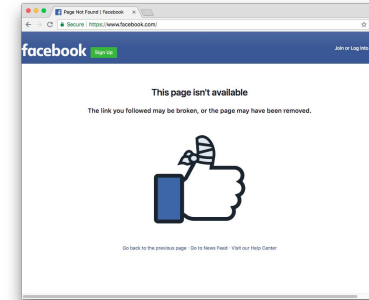
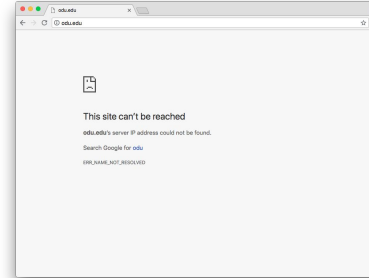
- Information access
- Personal interfacing (e.g., social media)
- Cultural artifact
- Personal knowledge reference



The Web is Ephemeral

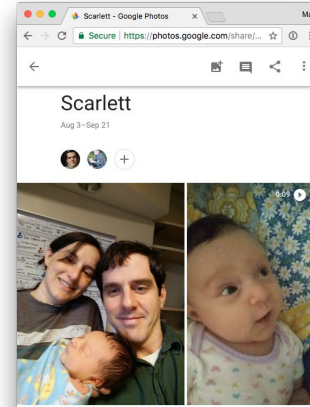
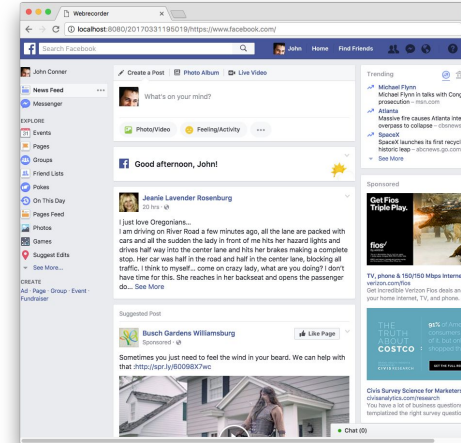


- No guarantee that information or services available today will be available tomorrow
- We are reliant on this service
 - Inclusive of mobile apps
- Links rot quickly
- User-generated content is especially fragile
 - Apps, social media posts, and comments vanish as services evolve



Our Web

- Born digital data
- Personalized Content
- Private Content
- Unique content
- Typically not preserved
 - Should it be?

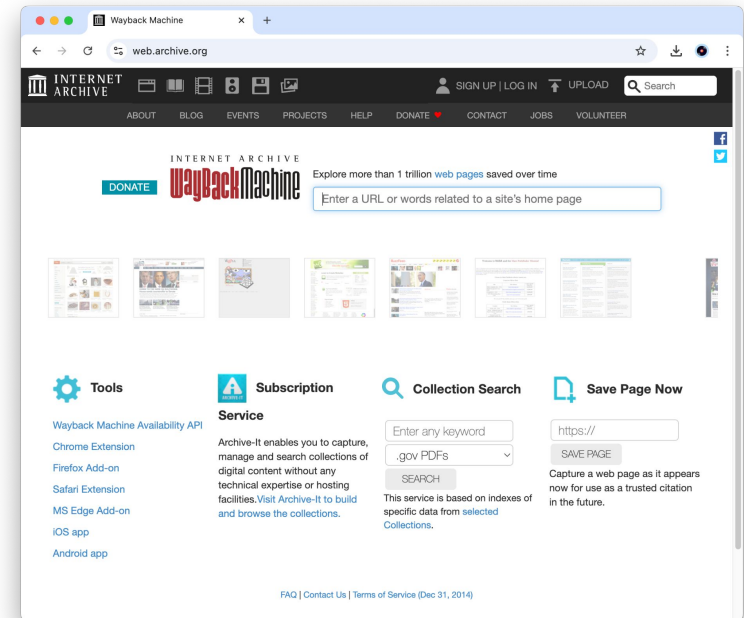


But Can We Save It?

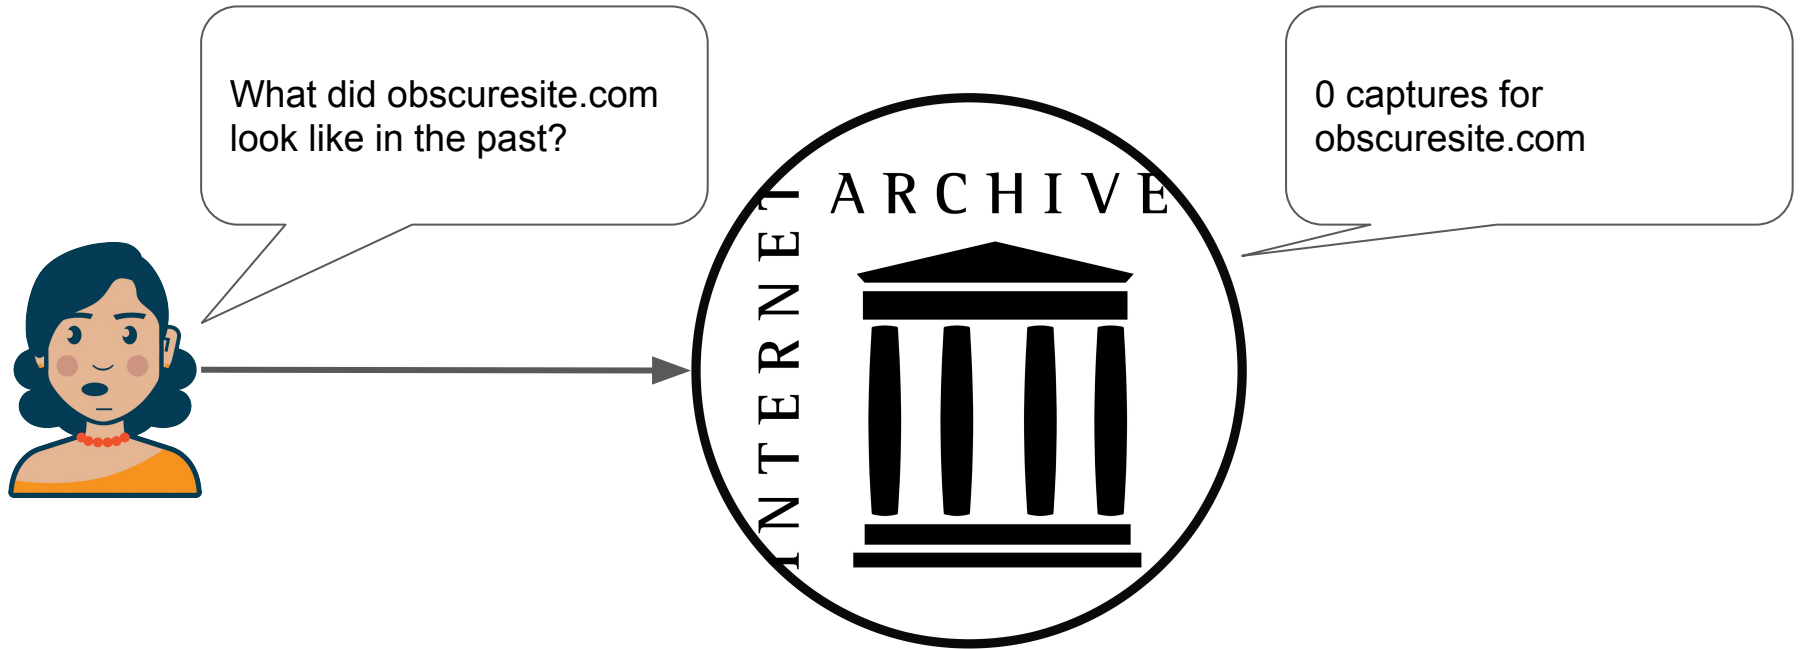
The Internet Archive

- Captures the web at 150 TB / day; currently 175 PB of web history (Dec 2025)
- Also: music recordings, radio broadcasts, podcasts, software, images, etc.

WayBackMachine



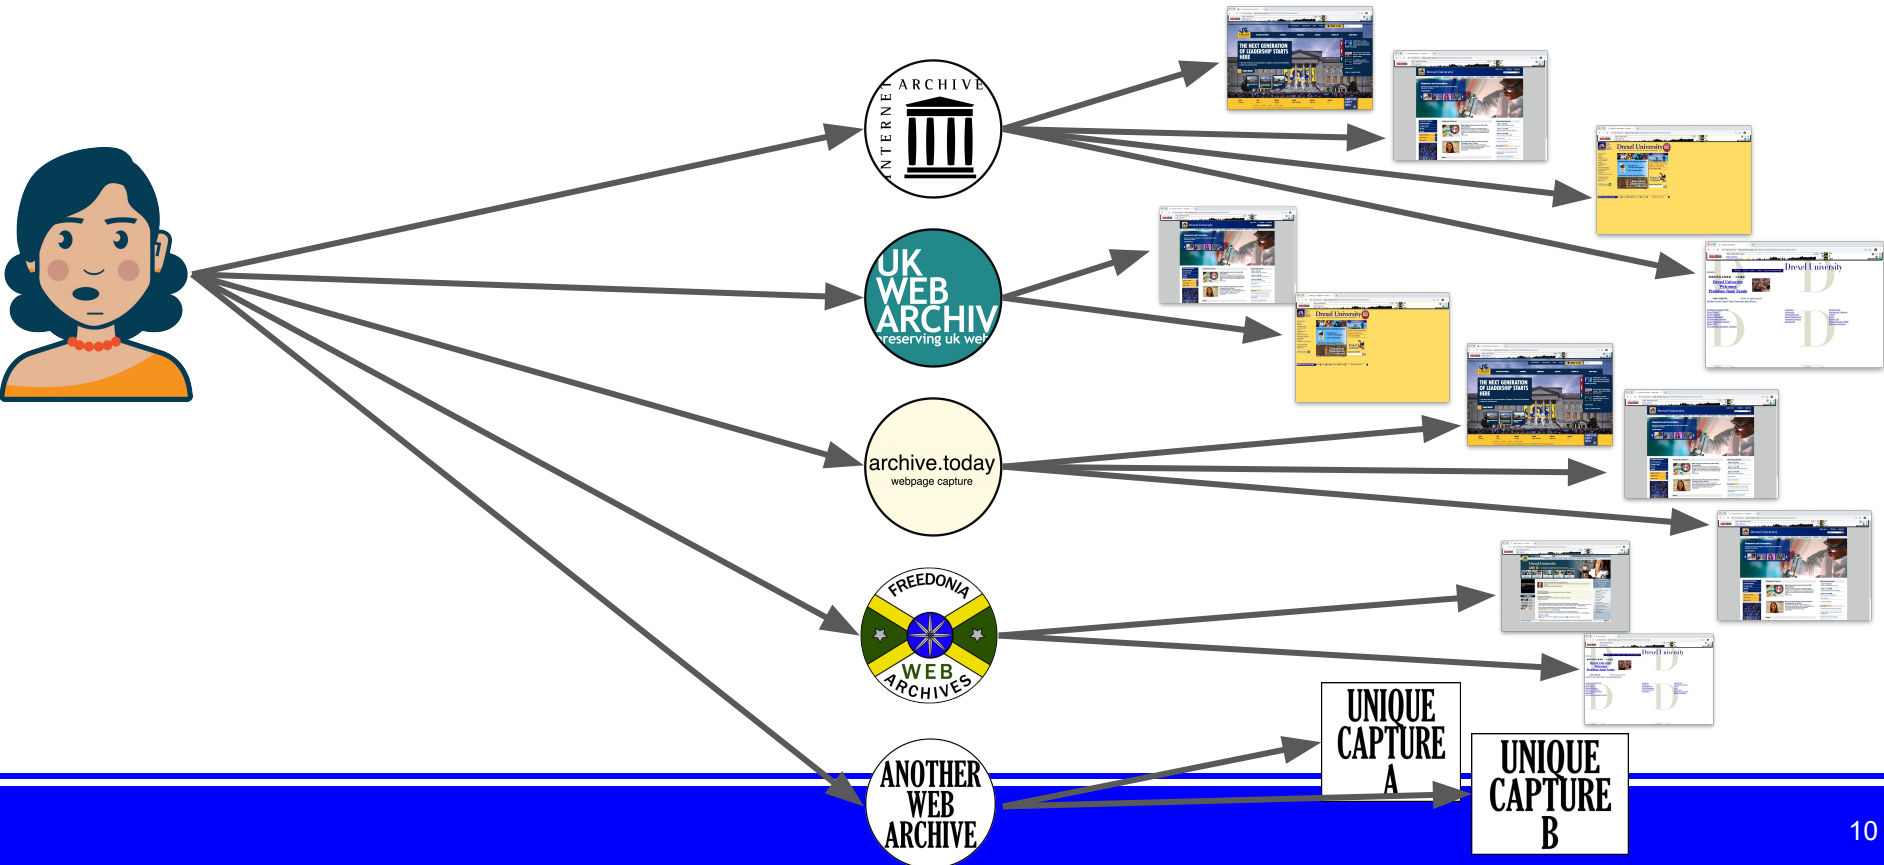
Not everything is preserved



Multiple archival efforts (3 of many)



More archives produces a more comprehensive picture



Web Archiving

Associate live Web URIs



With their archived representations

Web Archiving: not just “Save As...”

- Metadata
- Provenance
- Standards
- Big Data
- Composite Resources (cf. analog/conventional archiving)
- Technical Nuances
 - Indexing
 - Deduplication
 - Resilience
 - Handling Dynamicism
- Preservation is a race against time — the moment something appears online, it can begin disappearing

Web Archiving Metadata

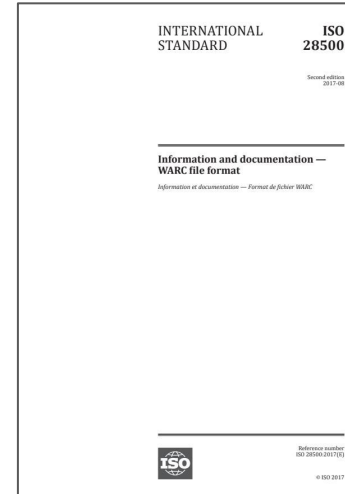
- The how, when and why web content was collected, preserved, and made accessible
- High-level descriptors
- Fields in crawls
- Derivative metadata formats
- Types:
 - Technical (format, checksums)
 - Descriptive (URL, pub date, site title)
 - Preservation (relationship w/ other captures, date/time, crawler used)
 - Administrative (curatorial decision, access restrictions)

Web Archiving Provenance

- Provenance documents where archived content came from, how it was collected, and when.
- Establishes the authenticity and integrity of archived web resources
 - Was it really as depicted? Who did it? When?
- Often realized through metadata
- Essential for reproducibility
 - Others can understand or repeat the capture process
- Reveals biases or gaps
 - what was included, excluded, or technically unreachable
- Clarifies chain of custody when content is transformed, migrated, or redacted over time

Web Archiving Standards

- WARC: ISO 28500:2017
 - Storage format for web archives
 - Container that contains transactional information and metadata
- Memento: IETF RFC7089
 - Provides mechanisms for interacting with the past web



Independent Submission
Request for Comments: 7889
Category: Informational
ISSN: 2070-1721

H. Van de Sompel
Los Alamos National Laboratory
M. Nelson
Old Dominion University
R. Sanderson
Los Alamos National Laboratory
December 2013

HTTP Framework for Time-Based Access to Resource States -- Memento

Abstract

The HTTP-based Memento framework bridges the present and past Web. It facilitates obtaining representations of prior states of a given resource by introducing datetime negotiation and TimeMaps. Datetime negotiation is a variation on content negotiation that leverages the given resource's URL and a user agent's preferred datetime. TimeMaps are lists that enumerate URIs of resources that encapsulate prior states of the given resource. The framework also facilitates recognizing a resource that encapsulates a frozen prior state of another resource.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see [Section 2 of RFC 5741](#).

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7889>.

Web Archiving Standards: WARC

- ISO-standardized file format (ISO 28500) used for storing web crawls and related digital content
- Uses records (e.g., response, request, metadata, revisit, resource) to represent different capture types
- Preserves original HTTP request/response pairs, supporting authenticity and replay
- Designed for long-term digital preservation: stable, extensible, and format-agnostic

WARC Anatomy

WARC response record

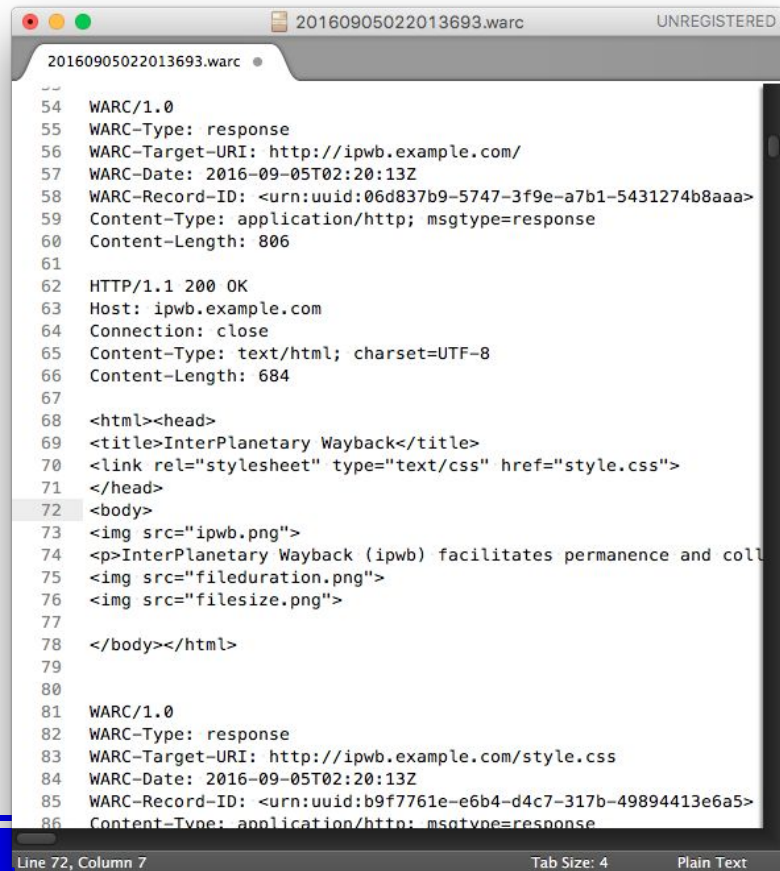
Warc-response
header

HTTP resp header

HTTP resp payload

WARCs also contain:

- HTTP requests
- warc-info
- warc-metadata records
- etc.



```
20160905022013693.warc
UNREGISTERED

20160905022013693.warc
54  WARC/1.0
55  WARC-Type: response
56  WARC-Target-URI: http://ipwb.example.com/
57  WARC-Date: 2016-09-05T02:20:13Z
58  WARC-Record-ID: <urn:uuid:06d837b9-5747-3f9e-a7b1-5431274b8aaa>
59  Content-Type: application/http; msgtype=response
60  Content-Length: 806
61
62  HTTP/1.1 200 OK
63  Host: ipwb.example.com
64  Connection: close
65  Content-Type: text/html; charset=UTF-8
66  Content-Length: 684
67
68  <html><head>
69  <title>InterPlanetary Wayback</title>
70  <link rel="stylesheet" type="text/css" href="style.css">
71  </head>
72  <body>
73  
74  <p>InterPlanetary Wayback (ipwb) facilitates permanence and coll
75  
76  
77
78  </body></html>
79
80
81  WARC/1.0
82  WARC-Type: response
83  WARC-Target-URI: http://ipwb.example.com/style.css
84  WARC-Date: 2016-09-05T02:20:13Z
85  WARC-Record-ID: <urn:uuid:b9f7761e-e6b4-d4c7-317b-49894413e6a5>
86  Content-Type: application/http; msgtype=response
```

Line 72, Column 7 Tab Size: 4 Plain Text

Web Archiving Standards: Memento



- Standardizes semantics and syntax for time on the web
- Content negotiation in the temporal dimension
- IETF RFC7089 (standard)



Original Resource

Resource that exists or used to exist;
we are interested in a past state of it

Link: `<URI-R>;type="original"`
points to Original Resource



Memento

Resource that encapsulates a past state of
the Original Resource

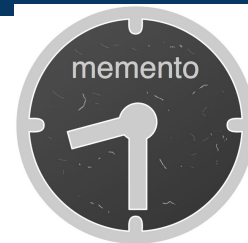
Link: `<URI-M>;type="memento"`
points to Memento



TimeGate

Resource that "decides", based on a given
datetime, which is the temporally best
Memento for an Original Resource

Link: `<URI-TG>;type="timegate"`
points to TimeGate



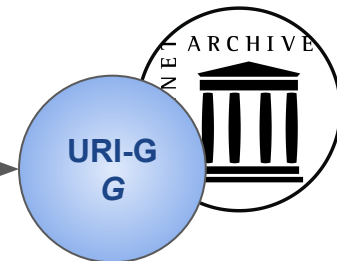
Background: Memento Request Example

HTTP Request

- **Accept-Datetime:** Wed, 02 Aug 2017 23:15:00 GMT
- **GET:** `http://web.archive.org/web/http://www.cnn.com`

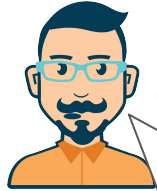


Request `cnn.com` at Sept
11, 2001 at 9am EST





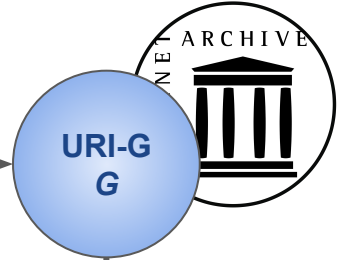
Background: Memento Request Example



HTTP Request

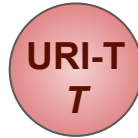
- **Accept-Datetime:** Wed, 02 Aug 2017 23:15:00 GMT
- **GET:** <http://web.archive.org/web/http://www.cnn.com>

Request `cnn.com` at Sept 11, 2001 at 9am EST



HTTP Response (302)

- **Memento-Datetime:** Wed, 02 Aug 2017 23:18:04 GMT
- **Location:** <http://web.archive.org/web/20170802231804/http://www.cnn.com/>
- **Link:**



timemap



original



timegate



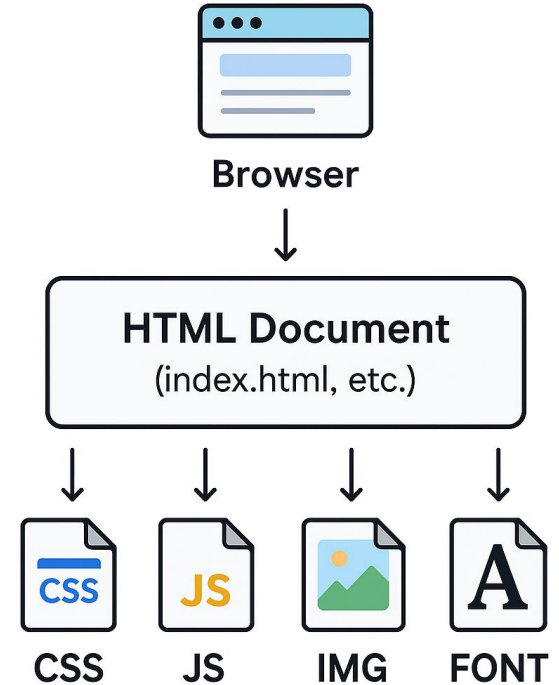
memento

Web Archiving (Big) Data

- Entire sites, documents, media, files
- Massive, complex, heterogeneous datasets, often measured in petabytes.
- Volume (billions of URLs harvested, longitudinal captures over the years)
- Variety (structured, semi-structured, unstructured content; media, scripts)
- Velocity (continuous capture of rapidly changing web content)

Composite Resources

- HTML → CSS, images, media → deeper embedded resources
- Content may appear dynamically based on user interaction

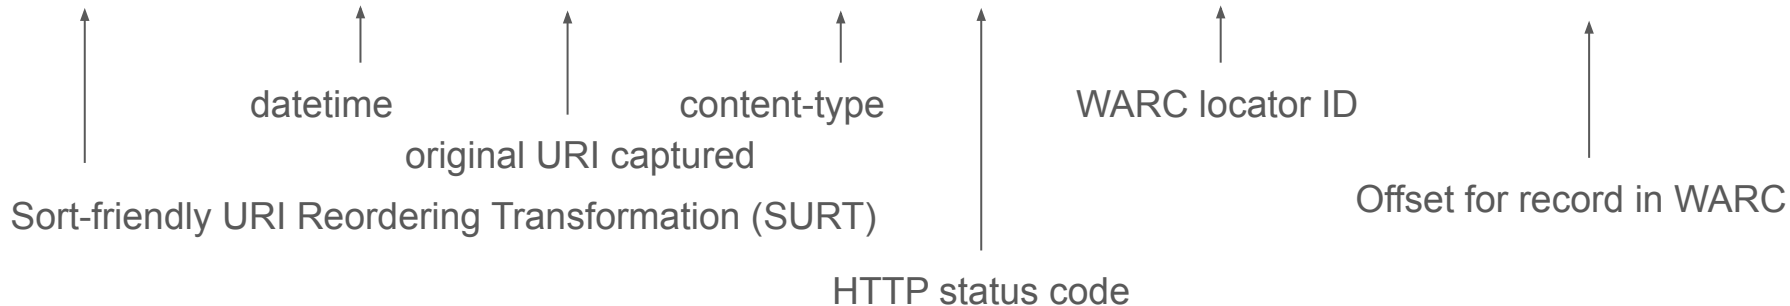


Indexing WARCs

- Fast lookup of a URL w/o scanning terabytes of WARC data
- Reply systems (e.g., Wayback, pywb) for reconstructing a page quickly
- Efficient queries over billions of archived documents
- Without indexing, random access is impossible
- Single WARC file can hold thousands of individual resources
- Index formats:
 - CDX (classic, well-defined fields)
 - CDXJ (modern, more flexible, supports extra metadata fields)
 - Database indexes (performance at scale, full-text search and complex queries)

Index Examples

```
edu,drexel)/ 19970626040823 http://www.drexel.edu:80/ text/html 200 NA4IS3YV4B0POE3KRG6VMNQLFFTZMWJG 1814
edu,drexel)/ 20220214060114 https://www.drexel.edu/ text/html 301 3I42H3S6NNFQ2MSVX7XZKYAYSCX5QBYJ 459
edu,drexel)/ 20230709040752 https://drexel.edu/ text/html 200 XMW7AUSKNSW02WMOML3TH2I3HVRVVHGD 30966
edu,drexel)/ 20251005193011 https://drexel.edu/ text/html 200 H5GRERCRQSZOWAEUV5APC7U67W2Z07LQ 32444
```



Deduplication

- Identify and handle duplicate or near-duplicate content
- Storage reduced, replay accurate, efficient indexing, avoid unnecessary redundancy
- Faster crawling, more efficient indexing, improved data quality with reduced noise, better preservation integrity
- Can be performed at crawl time or during replay
- Types:
 - URL-level (same page at different URLs)
 - Payload (identical content at different hosts)
 - Temporal (repeated captures unchanged over time)
 - Structural (template pages w/ minimal diffs)



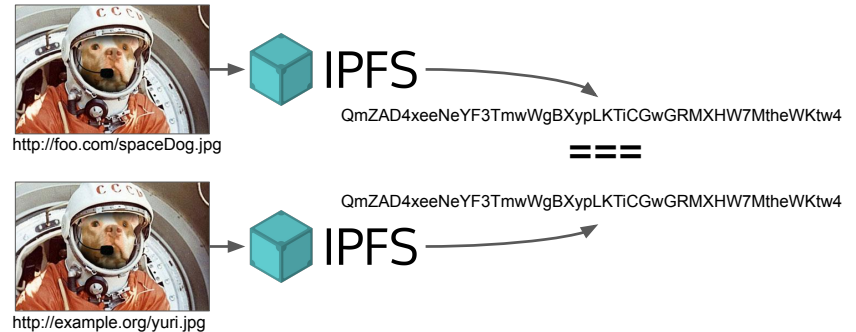
Screenshots of apple.com of the past

Resilience of Web Archives

- Wholesale or piecemeal copies of archives with means to check integrity
- How can you verify what is being purported as history is truly representative?
- Does authenticity matter if used for personal reference?



<https://github.com/oduwsdl/ipwb>



Handling Dynamicism

- The delta between a crawler and a browser causes resources to be missed
- JavaScript-driven, dynamic sites are hard to capture.

On the Change in Archivability of Websites Over Time

Mat Kelly, Justin F. Brunelle, Michele C. Weigle, and Michael L. Nelson

Old Dominion University, Department of Computer Science
Norfolk VA, 23529, USA
(mkelly, jbrunelle, mweigle, mln)@cs.odu.edu

Abstract. As web technologies evolve, web archivists work to keep up so that our digital history is preserved. Recent advances in web technologies have introduced client-side executed scripts that load data without a referential identifier or that require user interaction (e.g., content loading when the page has scrolled). These advances have made automating methods for capturing web pages more difficult. Because of the evolving schemes of publishing web pages along with the progressive capability of web preservation tools, the archivability of pages on the web has varied over time. In this paper we show that the archivability of a web page can be deduced from the type of page being archived, which aligns with that page's accessibility in respect to dynamic content. We show concrete examples of when these technologies were introduced by referencing memos of pages that have persisted through a long evolution of available technologies. Identifying these reasons for the inability of these web pages to be archived in the past in respect to accessibility serves as a guide for ensuring that content that has longevity is published using good practice methods that make it available for preservation.

Keywords: Web Archiving, Digital Preservation

1 Introduction

The web has gone through a gradient yet demarcated series of phases in which interactivity has become more fluid to the end-user. Early websites were static. Adoption of JavaScript allowed the components on a web page to respond to users' actions or be manipulated in ways that made the page more usable. Ajax [9] combines multiple web technologies to give web pages the ability to perform operations asynchronously. The adoption of Ajax by web developers facilitated the fluidity of user interaction on the web. Through each phase in the progression of the web, the ability to preserve the content deployed to the user has also evolved. As the web has evolved, the ability to capture and preserve the content as it is executed on the client side (i.e., within the browser after the page has loaded), it should follow that the archivability could be evaluated using a consistent replay medium. The medium used to archive (normally a web crawler tailored for archiving, e.g., Heritrix [21]) is frequently different from the medium used to replay the archive (henceforth, the *web browser*, the predominant means of

doi:10.1007/978-3-642-40501-3_5
doi:10.1007/s00799-015-0140-8

The Impact of JavaScript on archivability

Justin F. Brunelle · Mat Kelly · Michele C. Weigle · Michael L. Nelson

Received: 7 November 2013 / Revised: 12 January 2015 / Accepted: 14 January 2015 / Published online: 25 January 2015
© Springer-Verlag Berlin Heidelberg 2015

Abstract As web technologies evolve, web archivists work to adapt so that digital history is preserved. Recent advances in web technologies have introduced client-side executed scripts (Ajax) that, for example, load data without a change in top level Universal Resource Identifier (URI) or require user interaction (e.g., content loading via Ajax when the page has scrolled). These advances have made automating methods for capturing web pages more difficult. In an effort to understand why memos (archived versions of live resources) in today's archives vary in completeness and sometimes pull content from the live web, we present a study of web resources and archival tools. We used a collection of URIs shared over Twitter and a collection of URIs curated by Archive-It in our investigation. We created local archived versions of the URIs from the Twitter and Archive-It sets using WebCite, wget, and the Heritrix crawler. We found that only 4.2 % of the Twitter collection is perfectly archived by all of these tools, while 34.2 % of the Archive-It collection is perfectly archived. After studying the quality of these memos, we identified the practice of loading resources via JavaScript (Ajax) as the source of archival difficulty. Further, we show that resources are increasing their use of JavaScript to load embedded resources. By 2012, over half (54.5 %) of pages use JavaScript to load embedded resources. The number of embedded resources loaded via JavaScript has increased by

J. F. Brunelle (✉) · M. Kelly · M. C. Weigle · M. L. Nelson
Department of Computer Science, Old Dominion University,
Norfolk, VA 23529, USA
e-mail: jbrunelle@cs.odu.edu

M. C. Weigle
e-mail: mweigle@cs.odu.edu
M. L. Nelson
e-mail: mln@cs.odu.edu

12.0 % from 2005 to 2012. We also show that JavaScript is responsible for 33.2 % more missing resources in 2012 than in 2005. This shows that JavaScript is responsible for an increasing proportion of the embedded resources unsuccessfully loaded by memos. JavaScript is also responsible for 52.7 % of all missing embedded resources in our study.

Keywords Web architecture · Web archiving · Digital preservation

1 Introduction

How well can we archive the web? This is a question that is becoming increasingly important and more difficult to answer. Additionally, this question has significant impact on web users [40, 43] and commercial and government compliance [52, 53, 66].

The web has gone through a gradient of changes fueled by increasing user demand for interactivity. Early websites were relatively static, while continued adoption of web technologies has made the pages personalized and more interactive. JavaScript, which executes on the client, provides additional features for the web user, enabling or increasing interactivity, client-side state changes, and personalized representations. These additional features offer an enhanced browsing experience for the user.

As the web has evolved, the ability to capture and preserve the content as it is executed on the client side (i.e., within the browser after the page has loaded), it should follow that the archivability could be evaluated using a consistent replay medium. The medium used to archive (normally a web crawler tailored for archiving, e.g., Heritrix [21]) is frequently different from the medium used to replay the archive (henceforth, the *web browser*, the predominant means of

to perform synchronous client-server interactions after the HTML is loaded. The first wide-scale implementation of Ajax was in Google Maps in 2005, but Ajax was officially added as a standard in 2006 [70]. While archival tools per-

Springer

The Archival Acid Test: Evaluating Archive Performance on Advanced HTML and JavaScript

Mat Kelly, Michael L. Nelson, and Michele C. Weigle

Old Dominion University
Department of Computer Science
Norfolk, Virginia 23529 USA
(mkelly, mln, mweigle)@cs.odu.edu

ABSTRACT

When preserving web pages, archival crawlers sometimes produce a result that varies from what an end-user expects. To quantitatively evaluate the degree to which an archival crawler is capable of comprehensively reproducing a web page from the live web into the archives, the crawler's capabilities must be evaluated. In this paper, we propose a set of metrics to evaluate the capability of archival crawlers and other preservation tools using the Acid Test concept. For a variety of web preservation tools, we examine previous captures within web archives and note the features that produce incomplete or unexpected results. From there, we design the test to produce a quantitative measure of how well each tool performs its task.

Categories and Subject Descriptors

H3.7 [Online Information Services]: Digital Libraries and Archives

General Terms

Experimentation, Standardization, Verification

Keywords

Web crawler, Web Archiving, Digital Preservation

1. INTRODUCTION

Since much of our cultural discourse occurs on the web, web archiving is necessary for posterity. The goal of web archiving is to capture web pages so that they can be "replayed" at a later date. Web archiving tools access these pages on the live web in a manner similar to tools used by search engines (crawlers) and preserve the pages in a format that allows the data and contextual information about the page to be re-experienced. These "archival crawlers" take different approaches to digital preservation than their counterparts

Because archival crawlers attempt to duplicate what a user would see if he accessed the page on the live web, variance from what is preserved and what would have been seen compromises the integrity of the archive. The functional difference between archival crawlers and web browsers causes this sort of unavoidable discrepancy in the archives, but it is difficult to evaluate how good of a job the crawler did if the information no longer exists on the live web. By examining what sort of web content is inaccurately represented or missing from the web archives, it would be useful to evaluate the capability of archival crawlers (in respect to that of web browsers that implement the latest technologies) to determine what might be missing from their functional repertoire.

Web browsers exhibited this deviation between each other in the early days of Web Standards. A series of "Acid Tests" that implemented the Web Standards allowed each browser to visually and functionally render a web page and produce an evaluation of how well the browser conformed to the standards. In much the same way, we have created an "Archival Acid Test" to implement features of web browsers in a web page. While all standards-compliant browsers will correctly render the live page, this is not always the case when the archived version of the page is rendered. This difference can be used to highlight the features that archival crawlers are lacking compared to web browsers and thus emphasize the deviations that will occur in web archives compared to what a user would expect from a digitally preserved web page.

2. RELATED WORK

Web archives are generated by a variety of tools in a variety of formats. An ISO standard format utilized by institutional and personal web archivists alike is the Web ARChive (WARC) format [1]. WARC files allow HTTP communication that occurred during a crawl as well as payload, metadata and other archival features to be encoded in a single or an externally defined set of WARC files.

Heritrix paved the way for Internet Archive (IA) to utilize WARC files. IA's Wayback Machine (IA's Wayback Machine (and its open source counterpart¹) is the de facto standard.

Multiple services exist that allow users to submit URIs for preservation. IA recently began offering a "Save Page Now" feature co-located with their web archive browsing inter-

¹<https://github.com/ispac/openwayback>

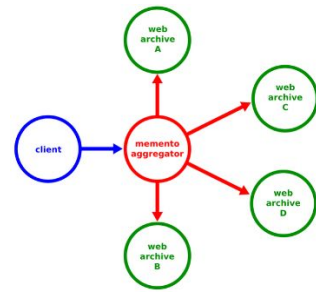
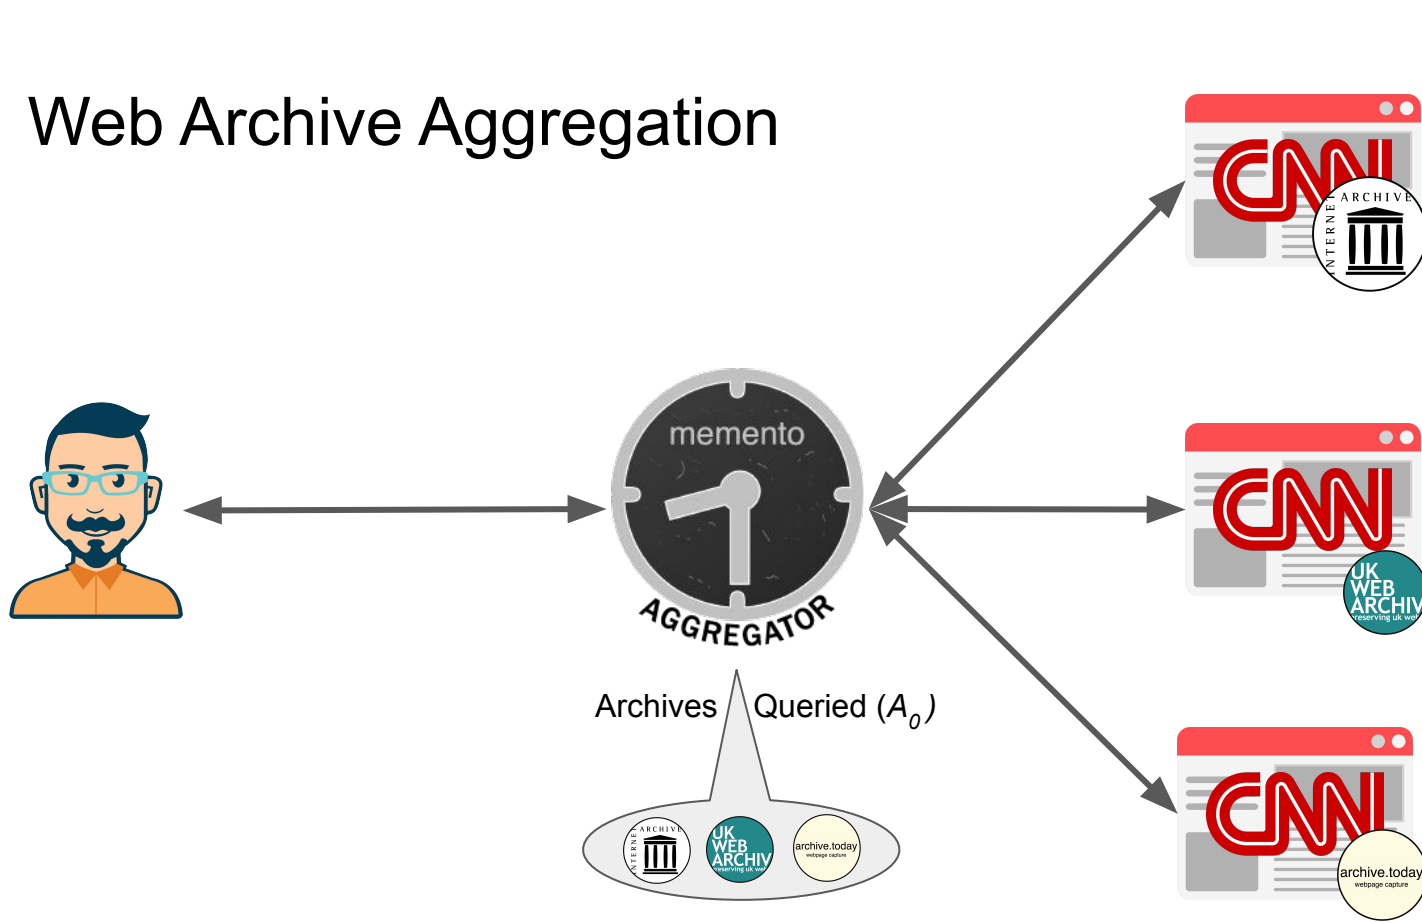
978-1-4799-5509-5/15/\$31.00 ©2014 IEEE.

doi:10.1007/978-3-642-40501-3_5

doi:10.1007/s00799-015-0140-8

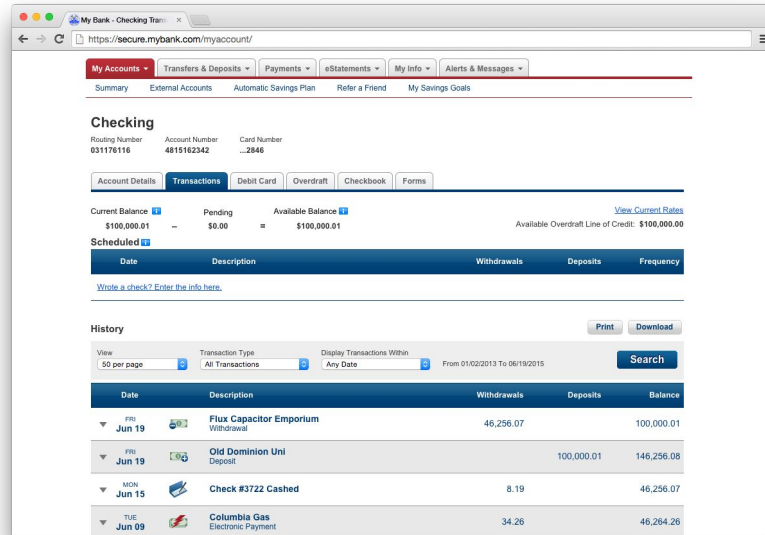
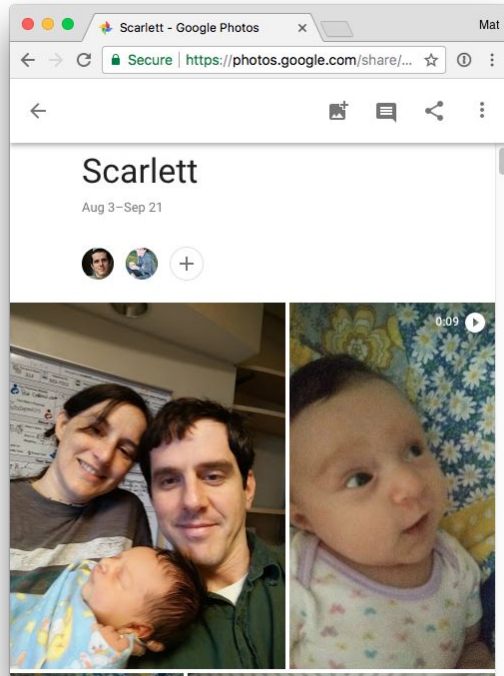
doi:10.1109/JCDL.2014.6970146

Web Archive Aggregation



Where Mat Fits

You Are Responsible for Saving What's Important



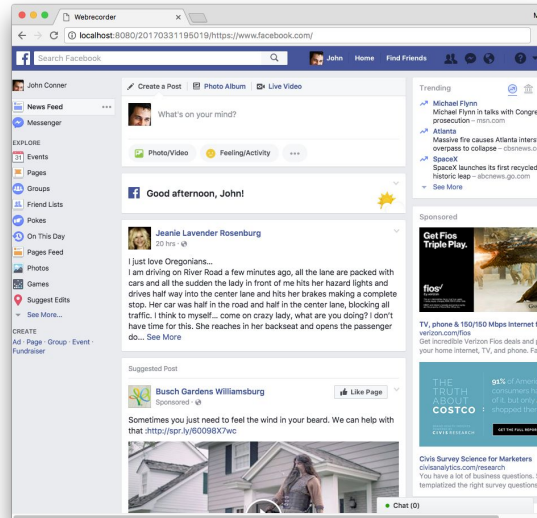
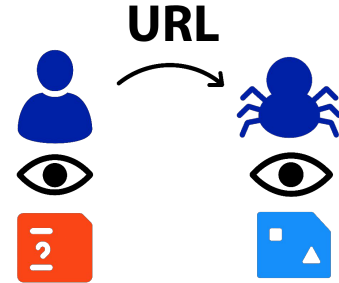
Most Web Archiving Tools Cater to Archiving the Public Web

- The face of society, Zeitgeist
- Easier to vet – captures from multiple institutions give legitimacy of the record
- Institutional self-interest
 - e.g., Stanford University Archives, UK Web Archive

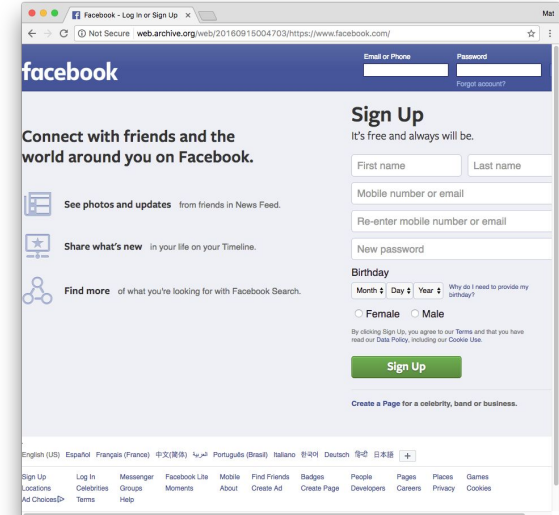


The Fallacious Key-Based Approach

- A URL is not enough
- Personalized Representations
- Dynamic web sites
- Context lost



facebook.com capture with browser-based archiving tools



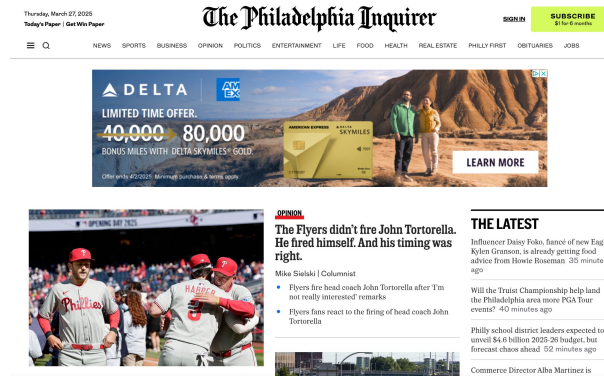
facebook.com capture from Internet Archive

Can We Save The Web We See From Our Perspective?

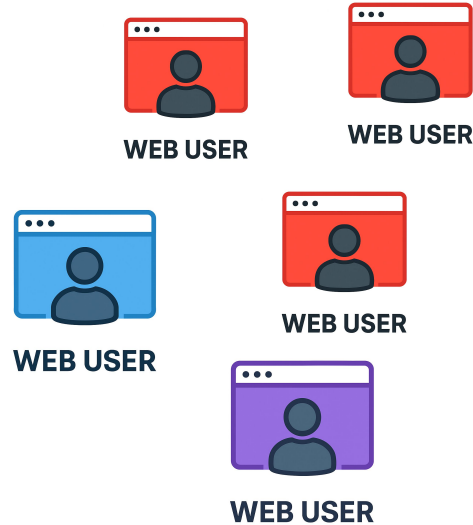
- Repurpose user's daily driver profile as crawler basis
- Permutate attributes of a user to represent a “persona”, producing a web experience closer to that of an actual web user cf. crawler
- Avoid clean slate crawling and delegation to a user-agnostic crawler
- Scale?



VS

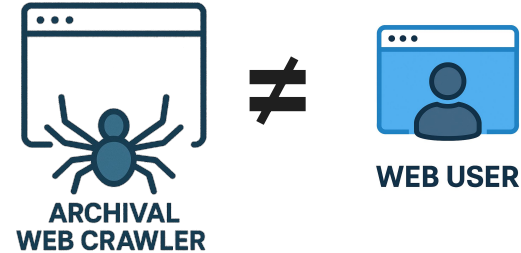


The Past Web Saved Is Not The Web That Was



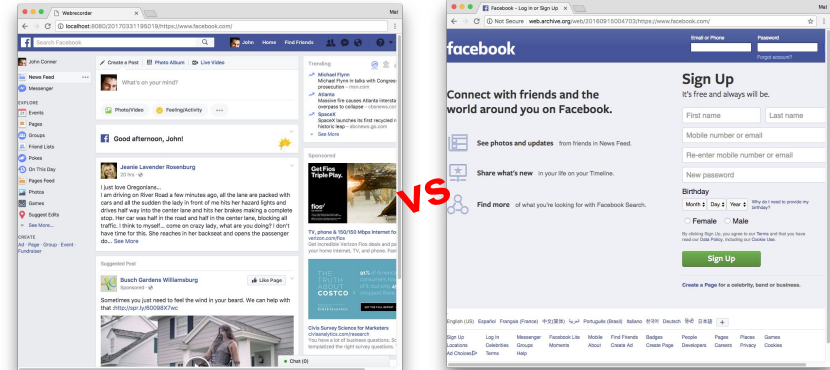
Crawlers Preserve A False Web (Premise)

- Archival crawlers preserve a version of the web inconsistent with web users' experience, a web that actually wasn't
- Customization, personalization based on user history is not *canonical*
- Crawlers (rightfully) see a clean/agnostic version of web sites, devoid of any individuals' experience, PII
- Ergo, what crawlers preserve is a version of the web inconsistent with what a user would have seen at that time
- False history? Nature of experience



By-Value Archiving with Posthoc Metadata Ascription

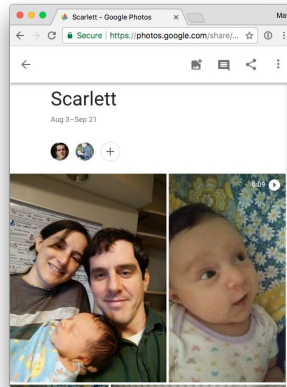
- Most true to the original form
- Capture complex content
- Personalized representations archivable
- Easier to fabricate content (not a good thing)
 - Requires a means of vetting authenticity



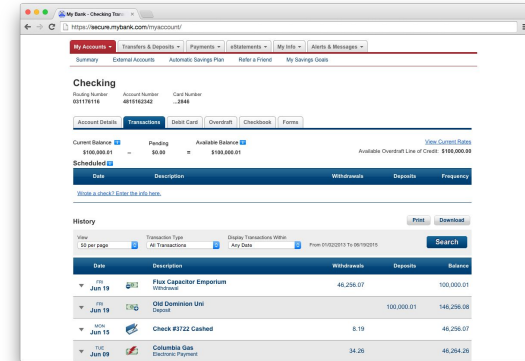
“Archive What I See Now: Bringing Institutional Web Archiving Tools to the Individual Researcher”, National Endowment for the Humanities Digital Humanities Implementation Grant (HK-50181-14)

Archiving Private and Personal Web Content

- The things we care about may not belong in an institutional web archiving
- Bearer of our own important information
- Stewardship, ownership, copyright
- Integration into conventional public web archive corpora

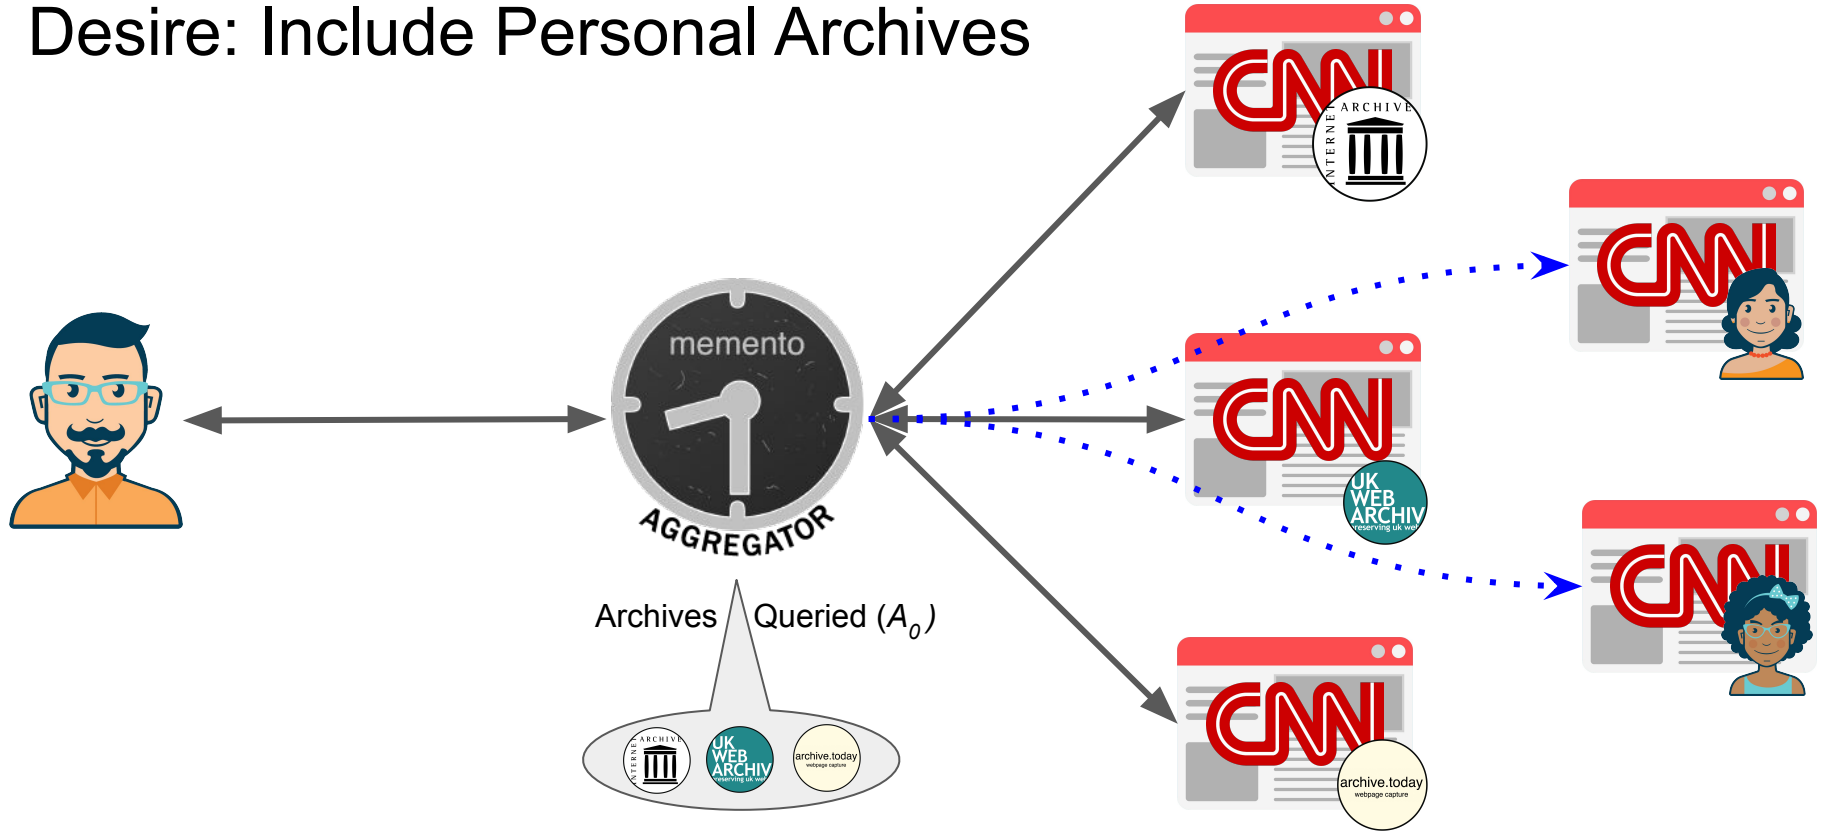


Historical capture of Google Photos

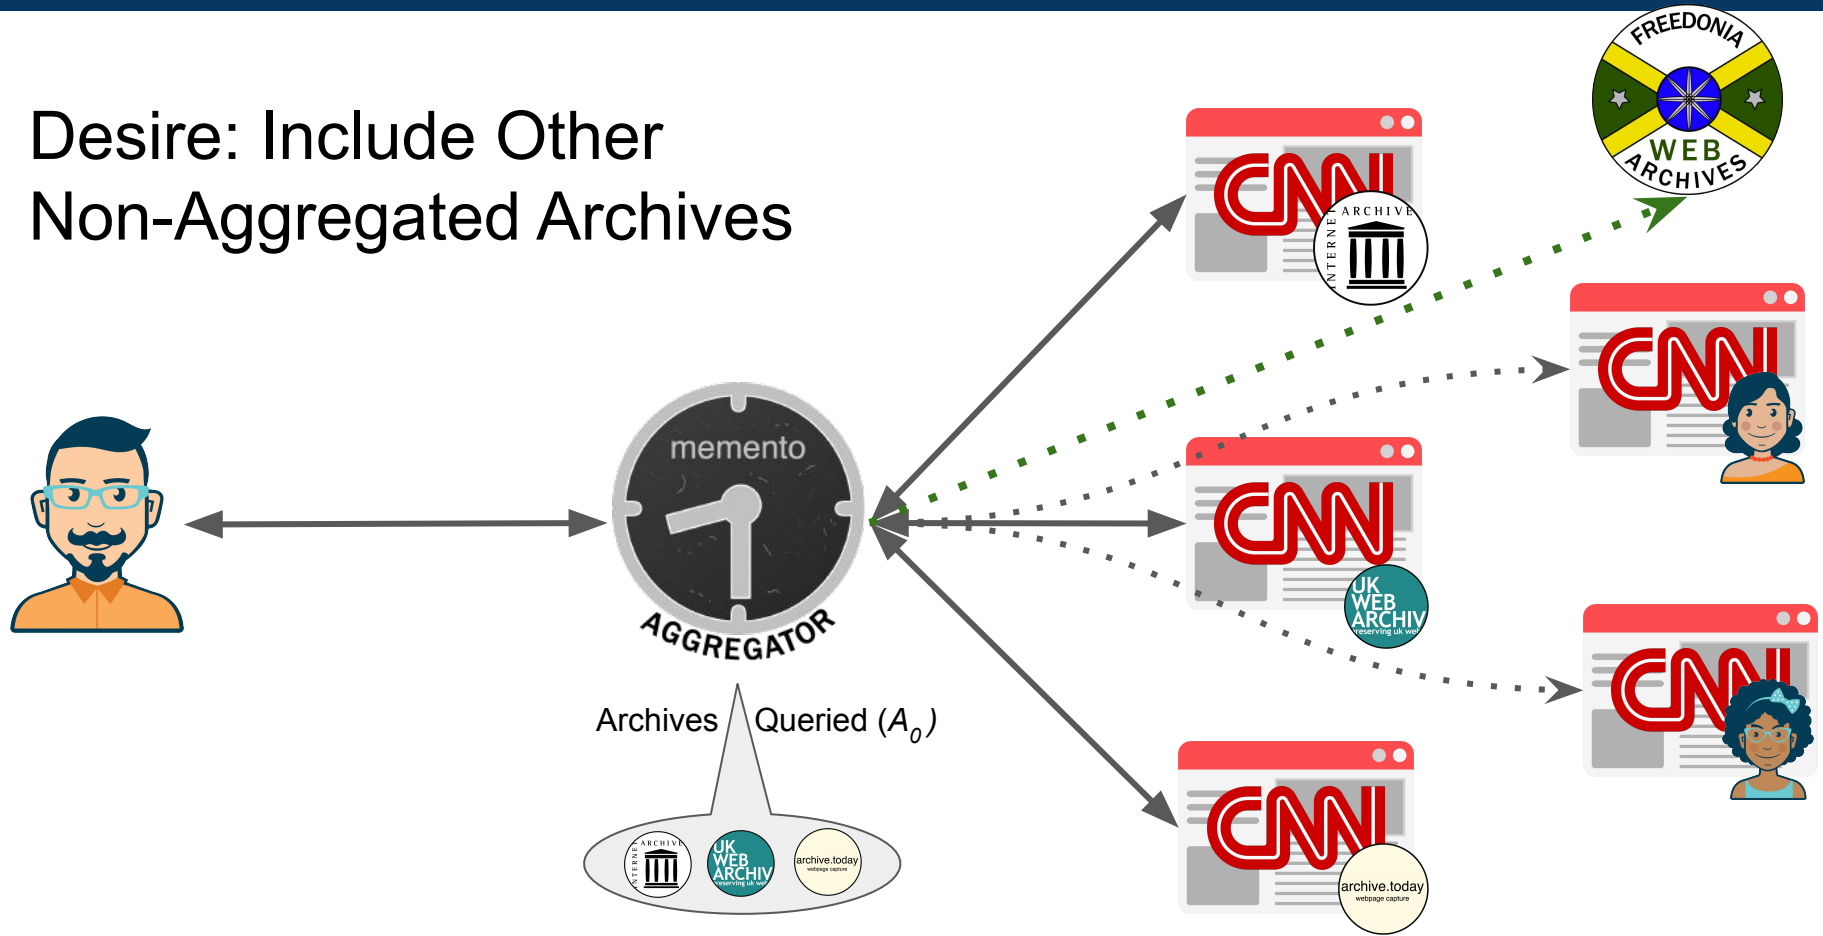


Historical capture of Online Banking Interface

Desire: Include Personal Archives

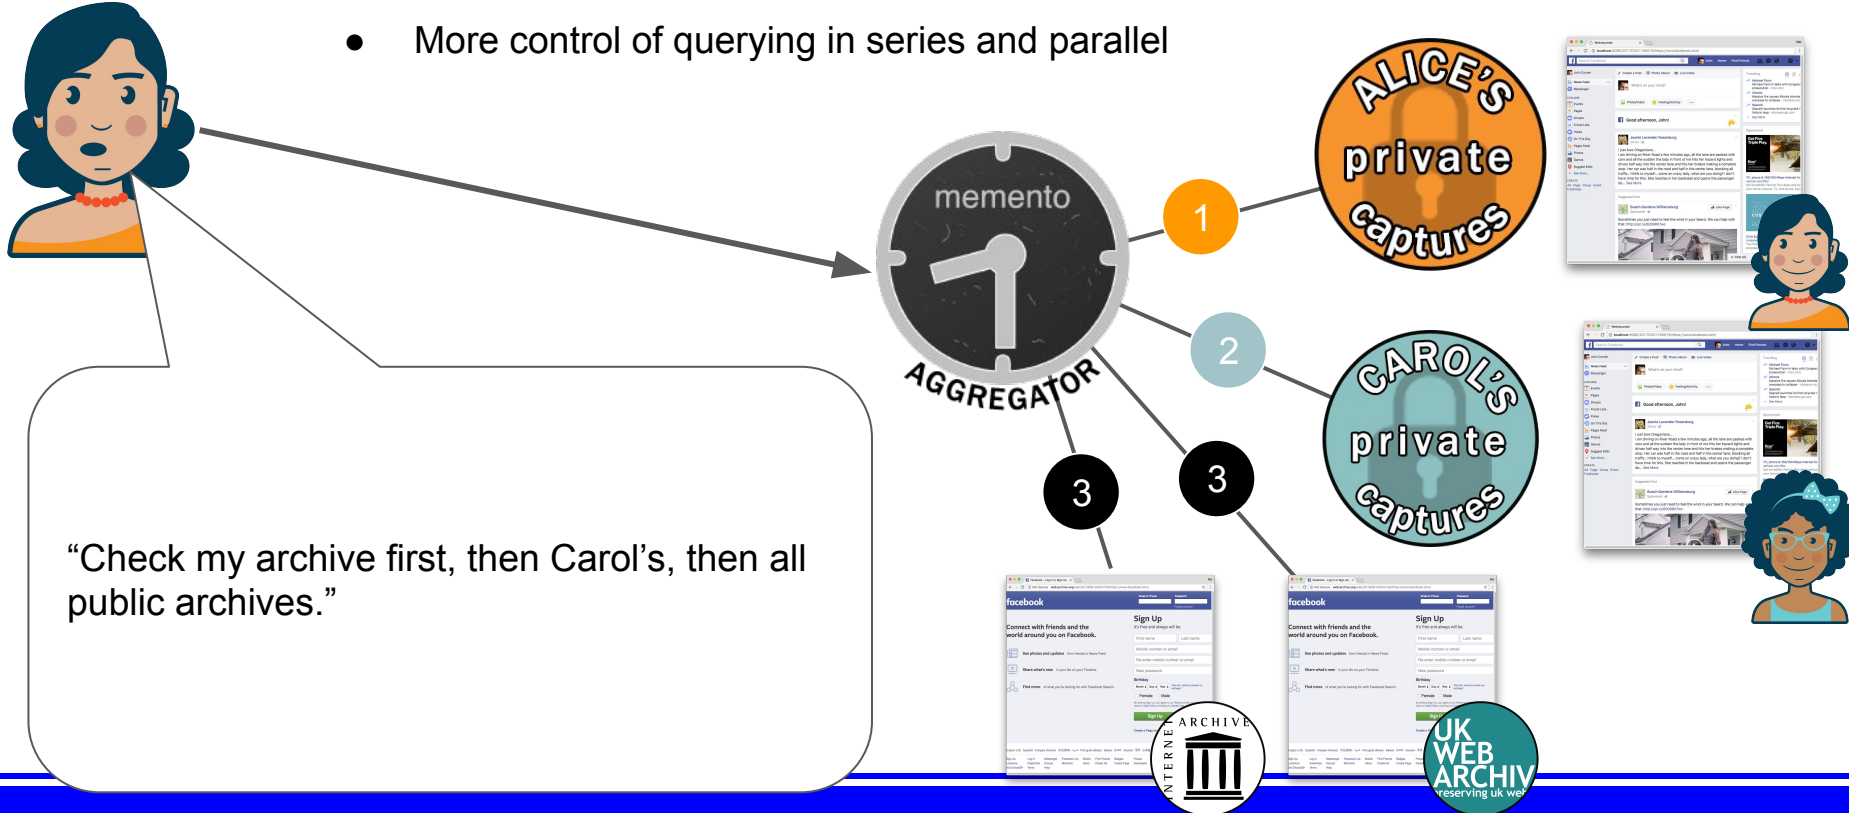


Desire: Include Other Non-Aggregated Archives



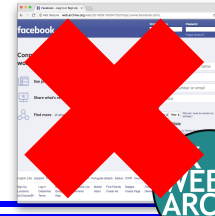
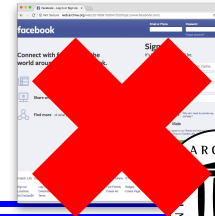
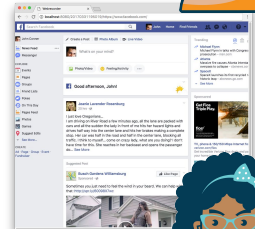
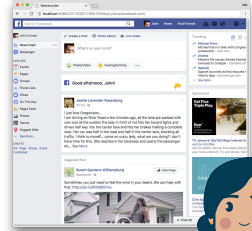
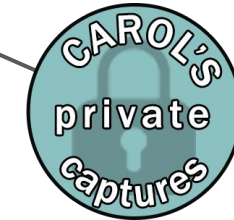
Query Precedence

- More control of querying in series and parallel



Query Short-Circuiting

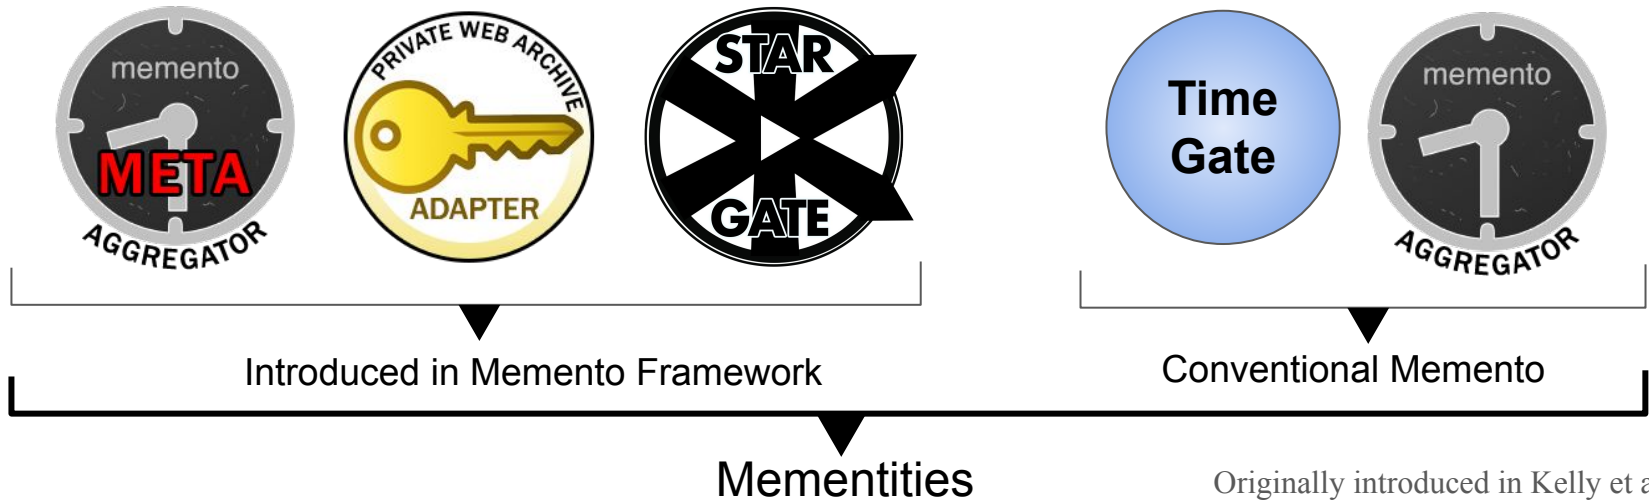
- May give priority to archive relevancy.
- **Series halt when threshold met.**



“Check private archives first. **Iff** you find no captures, only *then* check the public archives.

Mementities

- Memento + Entity (*entity* term already overused in web parlance)



Originally introduced in Kelly et al. 2018

Future Access

- Should these private photos/bank statements be accessible to anyone?
 - Strawman answer: no, but how can this be systematically implemented to be both sophisticated for security and privacy and usable by the non-technical
- How do we balance ability to access in the future with privacy and security?
- Should a page's public/private access now persist into the future?
- Reducing accessibility now can be at odds with any accessibility in the future

Can We Save The Web We See From Our Perspective?

- Repurpose user's daily driver profile as crawler basis
- Permutate attributes of a user to represent a “persona”, producing a web experience closer to that of an actual web user cf. crawler
- Avoid clean slate crawling and delegation to a user-agnostic crawler
- Scale?

Thursday, March 27, 2025
Today's Paper | Get Web Paper

The Philadelphia Inquirer

SIGN IN SUBSCRIBE \$1 for 6 months

NEWS SPORTS BUSINESS OPINION POLITICS ENTERTAINMENT LIFE FOOD HEALTH REAL ESTATE PHILLY FIRST OBITUARIES JOBS

PHILLIES
The Flyers didn't fire John Tortorella. He fired himself. And his timing was right.
Mike Siebeki | Columnist
• Flyers fire head coach John Tortorella after 1'm not really interested remarks
• Flyers fans react to the firing of head coach John Tortorella

THE LATEST
Influencer Daisy Foko, fiancé of new Eagle Kyleen Gramson, is already getting food advice from Howie Roseman 35 minutes ago
Will the Trust Championship help land the Philadelphia area more PGA Tour events? 41 minutes ago
Philly school district leaders expected to spend \$4.4 billion 2025-26 budget, but forecast chaos ahead 53 minutes ago
Commerce Director Alba Martinez is resigning from the Parker administration to produce a musical about Philly's Latino communities an hour ago
300-unit apartment building is planned for Conshohocken on SEPTA's land an hour ago
Dave McCormick defended DOGE and Trump, and called Signal chat a 'mistake' in first tele-town hall an hour ago
The Atlantic City aquarium was closed for 3 years but the 5th land one beloved wildlife exam on its back an hour ago
Butter up Philly! Things to Do 2 hours ago

PHILLIES opening day: Phils begin new season in Washington, pitcher traded to Marlins; roster includes three new key players
Updated 5 minutes ago
• MLB's experiencing outage issues on Opening Day
• Trea Turner, Kyle Schwarber to take turns hitting leadoff
• Ranger Suarez remains in Clearwater as he rehabs
• Ryan Howard recalls facing Zack Wheeler while he was still with the Mets

Take this Philly roads quiz and see if you can navigate like a local
Philly has nicknames for many of its highways and streets. Can you navigate the region based on how locals give directions?

VS

Thursday, March 27, 2025
Today's Paper | Get Web Paper

The Philadelphia Inquirer

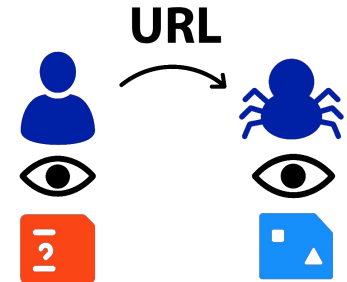
SIGN IN SUBSCRIBE \$1 for 6 months

NEWS SPORTS BUSINESS OPINION POLITICS ENTERTAINMENT LIFE FOOD HEALTH REAL ESTATE PHILLY FIRST OBITUARIES JOBS

DELTA
LIMITED TIME OFFER.
40,000 → 80,000
BONUS MILES WITH DELTA SKYMILES® GOLD.
Offer valid 4/20/25. Minimum purchase & taxes apply. [LEARN MORE](#)

PHILLIES
The Flyers didn't fire John Tortorella. He fired himself. And his timing was right.
Mike Siebeki | Columnist
• Flyers fire head coach John Tortorella after 1'm not really interested remarks
• Flyers fans react to the firing of head coach John Tortorella

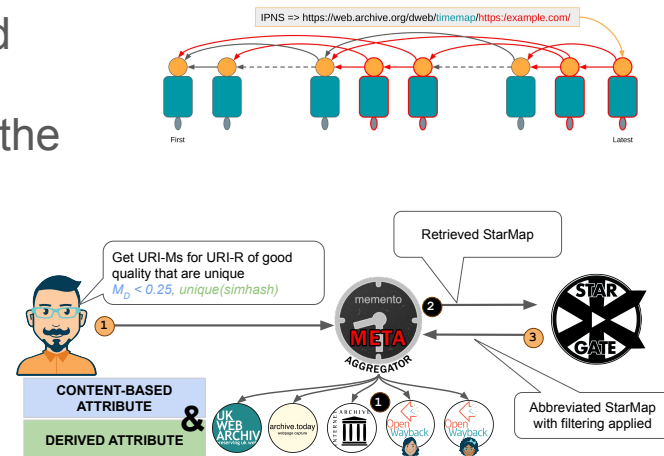
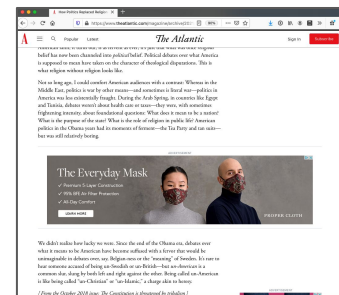
THE LATEST
Influencer Daisy Foko, fiancé of new Eagle Kyleen Gramson, is already getting food advice from Howie Roseman 35 minutes ago
Will the Trust Championship help land the Philadelphia area more PGA Tour events? 40 minutes ago
Philly school district leaders expected to spend \$4.4 billion 2025-26 budget, but forecast chaos ahead 52 minutes ago
Commerce Director Alba Martinez is



Outstanding Problems Under Investigation

A few of many

- Web Advertisements are intentionally lost but are important
 - AI-driven personas to surface deep web content; reusing browser profile to inform crawls
- Private/Personal web archive resilience through replication with security
- Resilience of remote resources through sophisticated linked data structures
- User-informed preference for information retrieval of the past in dimensions beyond time
 - e.g., topical, quality, ad hoc supplied lambda functions
- **Better** defining fuzziness of public/private and personal/institutional captures to better align with non-boolean reality



-30-