



# On Interdisciplinary Studies of Physical Information Infrastructure

*lewis levenberg*

Sometimes, in order to answer the questions that we ask as researchers, we need to combine more than one way of thinking about our questions. In situations when our objects of inquiry—the things in the world we question—are made up of some mixture of people, processes, and/or systems, we may find unusual juxtapositions of research methods particularly useful for discerning the important issues at stake. In these cases, our research benefits from the flexibility and breadth that we can bring to the ways that we ask, and answer, our research questions.

When we embrace this methodological diversity, we can decide how important it is for our approach to be **replicable**, our results **reproducible**, or our argument **intuitive**, and we can select a set of specific techniques and methods that fit these priorities. For example, to examine how **public policy** and **physical infrastructure** affect large-scale **digital communications**, we can use practices and techniques from any or all of computer science, policy analysis, literary studies, sociology, and history,

---

l. levenberg (✉)

Levenberg Services, Inc., Bloomingburg, NY, USA  
e-mail: work@lewislevenberg.com

© The Author(s) 2018

l. levenberg et al. (eds.), *Research Methods for the Digital Humanities*,  
[https://doi.org/10.1007/978-3-319-96713-4\\_2](https://doi.org/10.1007/978-3-319-96713-4_2)

in varying proportions. In this chapter, I will demonstrate, through examples, how a beginning researcher of **information infrastructure**—the physical and technical elements of how we move information around—might use a similar combination of methodological approaches.

To illustrate the broad variety of specific research techniques that a Digital-Humanities perspective makes available, I will introduce examples from my own research on the **network infrastructure** across West Africa, where large-scale **telecommunications network architecture** remains unevenly distributed, both between and within the region's nation-states. Growth rates of **backbone** network infrastructure (the internet's 'highways', large-scale connections between cities, countries, or continents) outpace growth rates of access to those same networks for people in this region. The pattern contrasts with most of the rest of the world's demand-driven network development (in which larger backbone elements are only constructed once demanded by the scale of the networks trying to connect to each other). The internet—comprised of **physical network** infrastructure, technical **protocols, software**, and the movement of **data**—commonly enters public **discourse** in terms of a borderless, international, global phenomenon. On this basis, we might expect that telecommunications **policies** to influence how internet infrastructure is developed would tend to come from global political powerhouses. Yet we find that Ghana, Nigeria, and even Liberia, despite their apparent relative weakness in a **geopolitical** context, appear to pursue aggressive telecommunications policy with broader, regional effects. This anomaly raises the central **research question** that I wanted to answer: How and why would the telecommunications policy strategies of these ostensibly weak states lead to a backbone-first **architecture** for large-scale internetworking throughout West Africa?

Because a researcher might approach such a question from any of various perspectives, we have an opportunity to separate the project into its **epistemological** and **methodological** pieces. Epistemologically, the focus on its object of analysis—the effects of these case studies' telecommunications policies on large-scale physical network infrastructure over time—gives the study its place astride technological and humanities ways of thinking. Methodologically, the approach to each piece of this puzzle highlights its own set of research practices. For example, to understand the actual changes to the region's network infrastructure over time, we can use scanning and **topological** analysis techniques from computer science. To find out how we know what a particular telecommunications

policy “means” or “intends”, we can use both **computational** and **heuristic** reading practices. And to articulate whether there was actually a causal relationship between telecommunications policies and network infrastructure changes, we can use higher-order analyses from historical and public-policy perspectives. By combining the results of these methods, the study arrives at a unified argument to answer its question: the backbone-first telecommunications policies of these ostensibly weak states are rational initiatives in their historical context, and they have disproportionately effective results on large-scale network architecture across the region. This is because their policies rely on, favor, and reinforce the states’ “**gatekeeper**” style **institutions of governance**—structures that work to concentrate political-economic **power** (and the perception of that power) at the state’s physical and conceptual boundaries.

### PROCEDURE: LAYERING METHODS

In holistic Digital Humanities studies of information infrastructure, we cannot rely solely on the selection of any given techniques from various disciplines. In addition to selecting our research methods pragmatically, for their relative efficacy at answering a part of a research question, we must also attend to the way in which those methods complement or contradict one another. In my study on West African network backbone infrastructure, I use the tools of different humanities, social-sciences, and computer science disciplines depending not only on the type of information that they help glean, but also on how they can build upon one another as I move through the phases of the study. Just as the architecture of information infrastructure includes discrete “**layers**” of machines, processes, human activity, and concepts, so too does the study of that architecture allow for multiple layers of **abstraction** and assumption, each a useful part of a unified, interdisciplinary approach.

To that end, I began my own study with background work, in the form of **historical** research. I reviewed the major developments in the cultural and political conditions of the region, and of each of the case studies, from 1965 to 2015. The challenge in this particular **historiography** was to connect the development of global internet networking, which mostly took place in North America and Europe at first, to the specific changes taking place in each of the West African case study countries during that period. In broad strokes, the social transformations in

Ghana, Nigeria, and Liberia in the late twentieth century were quite isolated from the technological work underway to create the internet. However, by concentrating on the **structural** forms of emerging **institutions**, we see in that period the first hints of gatekeeper-friendly governance, in both the case study states and for the internet.

Although that historical argument is superfluous to the present, methodologically focused chapter, I refer to it here in order to emphasize the importance of multi-disciplinary approaches to complex questions. Historical and regional-studies analysis of **macroscopic** narratives helped focus the next phases of my inquiry, by validating some **assumptions** about the time frame and relevance of my questions, and by providing necessary context to the examination of policy-making in the very recent past across those cases. In other words, it established the **historicity** of the material infrastructure under examination in the study. From there, I could undertake network data collection, and collection of the text **corpora**, with confidence that the period that I was studying was likely to prove itself significant for answering my question. These data-collection techniques would have worked just as well for studying any given period, but because I introduced them based on the historical layer, it helped validate some of my core **assumptions** before I delved too deeply into minutiae.

Next, based on the data I had collected, I moved into the **analytical** techniques, looking for the overarching patterns across the case studies. Without the data-collection stage for the network data, my network analysis would have been based on secondary sources. Likewise, without the text collection stage, I would have had to rely on the judgment of others to select which policies I would go on to read closely. And I was able to undertake my close reading of selected texts based on their prevalence in the **indexing** techniques I used across the whole collection. I could therefore lean on my own **interpretive** analysis of the texts in this phase, rather than on the biographical, historical, or other external context of the text or its creator.

I added other layers of abstraction as I collected more information from these disparate techniques, building towards the **theoretical** arguments that I would go on to make about the patterns that I perceived. Zooming out, as it were, to that theoretical level, I was able to test my assumption that nation-states do act, through their agents, to impact the shape of international networks under their purview, because I could

review the evidence that I had already layered, through the earlier methods, of how these nation-states undertake their telecommunications policy agendas.

At this higher-order layer, I could draw the links between the changes I had identified in the network architecture of the region, distinctions between each state's political-economic structures. Theoretical literature helped frame how each of these weak states interacts with other (state and non-state) internet builders, how each uses the processes of internet-building and the products of internet connectivity to represent themselves to their citizens and to the international community, and how the logic that underpins their surprising impacts on large-scale internet architecture, from conception to implementation, is generated and reinforced.

From this vantage, I was able to articulate how and why backbone-first patterns of internetworking across the region occur, what they portend for the nation-states under study, and what I could **generalize** about this development strategy. The cases of Ghana, Nigeria, and Liberia demonstrate that the rationales of weak states, as they work to affect network-building, depend on the results of previous and ongoing network changes. They also depend on the particular **political economy** of weak states—but such network changes may themselves present significant political-economic potential for transitioning away from gatekeeper-state models in the region.

In the rest of this chapter, I will focus less on that specific set of data and arguments, and more on the specific methods that I used. First I will review the collection and analysis of network data, then the collection and analysis of unstructured texts, and finally the use of higher-order theoretical techniques for drawing inferences and conclusions based on the combination of those collections and analyses.

## COLLECTING AND ANALYZING NETWORK INFRASTRUCTURE DATA

To dissect how networks have changed over time, we can use both active scans of existing computers in those networks, and passive collection of already-existing data in and about the same networks. To begin collecting my own data, I enumerated the **autonomous system** numbers (ASNs) and groups of **IP addresses** assigned by the internet **registries** for each case study state. To do this, I first copied the publicly available IP address and ASN assignment **database** from the public **FTP** server

maintained by AFRINIC (the Africa region’s registry).<sup>1</sup> Studies taking place in other regions of the world would use the appropriate regional registry for such databases. Next, using the publicly available ‘ip2location’ database of correspondences between geographical coordinates and IP addresses, I filtered this list of possible addresses by the geographical location of each case study.<sup>2</sup> These initial queries resulted in a list of groups of IP addresses, known as classless inter-domain **routing** (CIDR) blocks. The sizes of these blocks, limited by the registries, ranges from a single IP address, to as many as 16,777,216 addresses (thankfully, such addresses are listed in order). To run scans against each unique IP address in these sets, I wrote a small Python script, using the “netaddr” module,<sup>3</sup> to expand each CIDR block into a list of all the individual IP addresses contained therein. For ease of use, I built one list of addresses for each of my three case studies, and kept them in plain text format.

The resulting data was appropriately formatted for the actual active scanning of whether a given address had anything using it. To do this, I used ‘nmap’, which sends out and tracks the response of small **packets** of data on TCP, UDP, or ICMP ports to any number of addresses. With nmap, and a helper program similar to nmap called ‘masscan’, I tested whether, across any of the open **ports** on each of the IP addresses in each list, there existed a listening **service**. Here again, output from this step of the process became the input for the next step, in which I re-scanned those internet-facing hosts directly, using the ‘curl’, ‘traceroute’, ‘bgpdump’, and ‘tcpdump’ programs to get more information about each of those remote systems. Supplementary information in that database included the types of **servers** in use, common **response** data from the servers, and the most-reliable **paths** for data to travel to and from those endpoints. These tests resulted in my own ‘snapshot’ of current internet-facing infrastructure throughout the region. Repeating this process regularly for several months revealed clear growth patterns in that infrastructure over time.

To repeat this process, you need access to a computer with an internet connection. Each of the software programs and databases that I used are open-source and publicly available; I ended up needing to write very

<sup>1</sup>“Index of ftp://ftp.afrinic.net/pub/stats/afrinic/.”

<sup>2</sup>“IP Address to Identify Geolocation Information,” *IP2Location*, accessed October 15, 2016, <https://www.ip2location.com/>.

<sup>3</sup><https://github.com/drkjam/netaddr>.

little of my own code for this project. While I used a GNU/Linux operating system on the local computer, and the correspondingly packaged versions of each software program listed above, as well as a plain text editor and the command line / terminal emulator installed on that machine, there are a vast number of resources—command-line and graphical alike—for replicating the same results using Windows, Mac OS, or any other modern operating system. The important parts of this method are not necessarily the specific software or tools that one uses, but the effort required to learn how to use one’s tools effectively as well as the patience to perform the research.

In my own study, I supplemented those scan results with extensive information from existing public datasets, as well as from recent, similarly constructed studies. The latter served the additional function of validating the methods and results of those recent studies, although those had different research questions, and arrived at their approaches and conclusions separately. Secondary sources included data sets from universities and independent research institutions, dedicated research projects, and the regional internet registries.<sup>4</sup> Such programs collect quantifiable network data on a regular basis, using replicable methods that any beginning network researcher would be wise to try for themselves.

World Bank and IMF program data, such as the Africa Infrastructure Country Diagnostic (AICD) database, also provided some baseline, conservative estimates of telecommunications infrastructure projects in the region, though they tended to under-count both the contributions and the infrastructure of local programs and institutions.<sup>5</sup> Industry reports

<sup>4</sup>“Index of <ftp://ftp.afrinic.net/pub/stats/afrinic/>,” accessed August 11, 2014, <ftp://ftp.afrinic.net/pub/stats/afrinic/>; “The CAIDA AS Relationships Dataset,” accessed September 20, 2015, <http://www.caida.org/data/as-relationships/>; Y. Shavitt, E. Shir, and U. Weinsberg, “Near-Deterministic Inference of AS Relationships,” in *10th International Conference on Telecommunications, 2009. Con\_TEL 2009*, 2009, 191–198.

<sup>5</sup>“Homepage | Africa Infrastructure Knowledge Program,” accessed June 3, 2016, <http://infrastructureafrica.org/>; “Projects: West Africa Regional Communications Infrastructure Project—APL-1B | The World Bank.”; “World Development Report 2016: Digital Dividends”; “Projects: West Africa Regional Communications Infrastructure Project—APL-1B | The World Bank”; Kayisire and Wei, “ICT Adoption and Usage in Africa”; “Internet Users (per 100 People) | Data | Table”; “Connecting Africa: ICT Infrastructure Across the Continent”; World Bank, “Information & Communications Technologies”; The World Bank, Financing Information and Communication Infrastructure Needs in the Developing World. Public and Private Roles. World Bank Working Paper No. 65.

and other third-party sources, particularly those provided for marketing purposes, provided additional, but less-easily verifiable, estimates of existing infrastructure and internet usage. For example, the Internet World Statistics website, or the Miniwatts Marketing Group report, reliably over-estimate technological data in stronger consumer markets, and underestimate the same phenomena in areas with less per-capita purchasing power.<sup>6</sup>

Together, these primary and secondary data sets outlined the general patterns of networking growth across the region, and provided more insight into network infrastructure and ownership in Ghana, Liberia, and Nigeria than had previously existed. Despite rapid infrastructure development, the network topologies of Nigeria, Ghana, and Liberia remain thin within each state. They are each becoming denser at the backbone tier, and along network edges, while last-mile infrastructure, hosting services, and internal networking are still lacking in the region.

These results illustrated the validity of my research question, by showing that weak states are indeed developing strong backbone networks. More importantly, they provided a set of clear, verified network outcomes against which I could benchmark the perceived impact of weak states' telecommunications policies. However, to do so, we must turn to a different set of research methods—the collection and analysis of textual data.

## COLLECTING AND ANALYZING UNSTRUCTURED TEXTS

To reliably collect and organize a large set of texts, we can use computational techniques from the branch of study known variously as **natural language processing**, **corpus linguistics**, **distant reading**, or **broad reading**.<sup>7</sup> This set of methods deals with the **aggregation** and analysis of large sets of **structured or unstructured texts**; different disciplinary approaches use distinct sets of algorithmic or programmatic approaches to understanding the contents of texts, depending on the purposes of their research.

<sup>6</sup>“Africa Internet Stats Users Telecoms and Population Statistics,” accessed December 16, 2014, <http://www.internetworldstats.com/africa.htm>.

<sup>7</sup>For simplicity, I refer to the specific techniques that I outline here by the latter term, but you can find excellent resources on these techniques, and their histories, under any of those names.

In the approach that I describe here, the purpose of this stage of research was to collect and organize a great deal of writing in a consistent **format**, and then to identify the most promising documents out of that collection for closer analysis. In other projects, researchers might be looking to identify the occurrence of specific, predetermined terms, or particular uses of language, across a large body of writing (as an aside, this colloquialism of a “body” of work is why the terms “corpus” and “corpora” are used in these disciplines), or to create maps of relationships between distinct texts in a corpus based on their **metadata**. In those cases, various natural language processing techniques certainly suit the need, but the specific work to be done would differ from what I describe here.

For my own broad reading of the early twenty-first century telecommunications policies of Ghana, Liberia, and Nigeria, I first collected about two thousand documents, web pages, and transcripts, each of which was produced by the government or the officials of these states. I collected public record and news archive searches, conference proceedings, policy-making negotiations, technical documentation, and legal documentation, mostly through the web interfaces of governmental and agency websites, libraries, and databases maintained by third-party institutions such as the World Bank. These texts came in a huge variety of **formats**, so one major, early challenge in cleaning this “corpus” was to convert as many as possible of the files I had copied into a **machine-readable** format. I used the open PDF **standard** for this purpose, although certain documents were rendered as images rather than texts, which required some manual intervention to apply **optical character recognition** settings and/or quickly scan the documents into a different format. As with so many other projects that include the collection of **unstructured** data (textual or otherwise), a great deal of effort and time had to go into the cleaning and preparation of the data before it was feasible to process using automated tools.

Once that was done, though, I was able to proceed to the **indexing** of these texts. I used the open-source program Recoll, which is based on the Xapian search engine library, to create an index of all the documents in my collection.<sup>8</sup> There are a vast number of such programs and products available, so the specific program or software library that you might

<sup>8</sup><https://www.lesbonscomptes.com/recoll/>.

use would certainly depend on your own project's requirements. At this stage, I used the index to search through all the texts at once, looking for higher-frequency occurrences of unusual words or phrases, such as proper nouns and technical terms.

Some of these patterns of language also helped classify the types of texts. For example, a frequent occurrence of terms related to the apparatuses, internal operations, and techniques of governance of the state, such as parliamentary names or legal references, tended to indicate texts produced by governmental bodies, such as legislative documents. Looser policy guidance documents tended to include more terminology distinct from the workings of the state itself, such as ideological or economic terms, or the use of terms like borders, security, identity, or control. Geographical or infrastructural terms were more likely than the baseline to appear in technical documentation, while historical or philosophical terms were likelier to pop up in speeches and arguments made by politicians and other public figures during policy negotiations or in transcripts.

The results of this process let me select a much smaller set of the collection of texts for my closer analysis, confident in the likelihood that each text would prove particularly technically, legally, academically, or politically significant. Next, I briefly reviewed each document in the selected set to ensure that I had eliminated duplicates, opinion articles and other similarly superficial pieces of writing, texts from countries or time periods outside the scope of my study, and other such false positives. Finally, I used the same indexing program from earlier, Recoll, to search through the full text of each selected document for terms that I expected to indicate the concepts that I would be reading closely for, such as “network,” “infrastructure,” and so on. Your own such keyword searches, when you get to this stage, will depend on the specific object and question that you are studying, of course.

At this juncture, a researcher will have selected just a few of the many documents in their collections, but they will be able to reasonably argue that these texts are likely to have the most relevant information for answering their own research question. Combined with an analysis of network infrastructure, we can now cross-reference metadata such as comparing the dates at which selected texts were published and disseminated with the dates when specific changes were measured in an observed network. However, we will not yet have understood the contents of these key documents. To proceed, we must do what humanities scholars have done for centuries—that is, we must read the texts for ourselves.

Any particular approach you take to reading a text closely will depend in large part on the specific research question you are trying to answer—and, by extension, the disciplinary tools at your disposal for answering it through reading. For example, in a literary study, you might read a text with special attention to its narrative structure, literary devices, or poetics. In a study of scientific discourse, you might instead read texts with a focus on the specialized terms, methods, results, or concepts that they cite, define, or change. And in a study of public policy outcomes, such as the one in this chapter's example study, close reading concentrates on identifying the ideological basis and implementation details of the policies described by the texts. As you can imagine, these different priorities can lead to very different observations and insights about the texts that you study.

In my own study's close reading, I hoped to understand how ostensibly weak states implemented large-scale technological policy, and effected very large-scale network changes. To that end, as I read through each text that I had selected, I noted the structural forms of the arguments that each policy set out as justification or incentive for its prescriptions, as well as the specific data that each policy cited. The overarching narratives and justifications for policy recommendations were not always explicitly named by the policies themselves. Still, enough of their underlying logic was expounded that when I would broaden my focus again to the overall theoretical and historical stakes of these policies' implementations, the close reading would serve as solid foundations for those arguments. I will not burden this chapter with examples of my close readings, but I mention the specific elements that I looked for because this was the part of the study that gathered my evidence for claims about how and why specific policies were enacted.

The basic techniques of close reading will be familiar to nearly all humanities researchers, but since they strike such a contrast to the computational research techniques that we have seen so far, we can spare a thought here for their mechanics. Effective close reading relies on your concentration on a specific text for a sustained period of time; speed-reading and skimming are useful techniques for gathering superficial information in greater volume from more texts, but they do not serve your ability to answer a research question through the observation of detail quite so well. Some useful tools in this context are quite ancient, but modern software for bibliographic management or for annotations of digital texts can support your readings. Whatever tools you prefer,

the crucial undertaking is to read the texts with attention to the details that you need to find in order to answer your research question—in these cases, to move from the details of telecommunications policy back out towards an historical, political analysis of information infrastructure.

### ARTICULATING THEORETICAL INSIGHTS

In this section, I connect the results of these disparate methods and disciplinary approaches to the larger project of ‘doing Digital Humanities’ in the study of information infrastructure. This is the point at which the evidence you have collected through each layered technique can be abstracted into an overarching argument, just as we rely on the underlying infrastructure of the internet to facilitate conversation with people around the world.

In my own study, I had found that, despite their noticeable political-economic and structural differences, Ghana, Liberia, and Nigeria each produced telecommunications network development policy specifically designed to target backbone network infrastructure development. Moreover, in each of these three cases, such policies have consistently preceded network infrastructure implementations which followed the architectures promoted by those policies. The network architectures of backbone-first or edge-first infrastructures are globally rare, and we would not expect to find many such architectural patterns in networks that are cobbled together ad hoc from existing smaller networks, or in the development of large centrally managed organizations. So, the fact that we find backbone-first architecture arising in the region where weak states have consistently produced technological policy calling for just such networks leads to the reasonable claim that these policies are working as intended.

Having examined how telecommunications policy affects ICT development beyond the expected capabilities of gatekeeper states in West Africa, I was then able to turn to the implications of that impact. Specifically, if network-building activity responds to aggressive telecommunications policy, then we can also observe long-term effects of network-building on the underlying political-economic conditions of the very states that produce telecommunications policies in the first place. Further, we might see here how weak, gatekeeping states can benefit from intensified investments of capital, labor, and policy attention on backbone infrastructure. Through the development of

telecommunications infrastructure and social institutions that rely on it, Ghana, Nigeria, and Liberia engage in a “thickening” of both physical and figurative networks. By reinforcing their positions as gatekeeper states, they lay the foundations of a transition away from that precarious structure, and towards more administrative governmental institutions. This argument stems from observing the implementation of the disproportionately aggressive policies of the case study states, and noting that network backbone infrastructure development correspondingly increased.

Other factors beyond the policies themselves also mattered. Improved funding and, in some cases, easier access to domestic and international capital for large-scale infrastructure investors, including telecommunications providers, also proved consistent during this period. At the same time, few technical policy alternatives to the backbone-first approach to telecommunications development were offered by non-state actors such as corporations or international institutions, in the documented policy discussions from any of the three case studies. Under these conditions, supposedly weak states can wield outsize influence on international network architecture. This observation then forms our criteria for inferring some degree of causal relationship described by the correlation between the telecommunications policy of our case studies, and the subsequent appearance of networks described by those policies. However, confounding factors in these policies’ implementations have ranged from funding sources for network ownership and construction that are located outside of the states in question, to economic demand for edge network providers (as opposed to local/end user/last mile demand). We can see that these confounding factors are accounted for by the weak states’ policy initiatives, which in turn allows us to acknowledge the awareness of each state of their weak position in geopolitical context.

In other words, Ghana, Liberia, and Nigeria have pursued policies that reinforce their existing political-economic structures. Gatekeeper states pursue gate-keeping internet infrastructure—the backbone networks and connections described throughout this study—as opposed to last-mile service provisions, or dense data storage centers (or the non-ICT infrastructure that would support those other ICTs). This feedback loop reinforces the leverage held by gatekeeper states during negotiations with external interests such as other governments or international bodies, because it increases the degree of control that these states have over

the “right of way” for global information flows. In this light, the backbone-first telecommunications policies of these ostensibly weak states do make rational sense, and they have disproportionately effective results on large-scale network architecture across the region, because the policies rely on, favor, and reinforce the states’ gate-keeping institutions of governance.

## CONCLUSION

In this chapter, I have introduced the wide variety of techniques available to researchers who study information infrastructure. The appropriation of specific approaches for discrete parts of a larger research project has its drawbacks, of course. We must be cautious not to cherry-pick our data or methods, selecting only those approaches or pieces of information that support our biases or keep us in our comfort zones. To that end, it is worth remembering that other methods than those outlined in this chapter can also be useful for studying physical information infrastructure. For example, conducting interviews with those who create and maintain infrastructure, or with policymakers, or different demographic segments of the networks’ users, would provide deeper insight into the perspectives of individual people on the issues at stake. Conversely, more targeted or sustained scans of specific hosts across the networks could provide further details of technological implementation of networking across the region, such as the distribution of routing or server software, or the relative usage of network address capabilities. The key to any particular combination of methods, however, remains its utility in the answering of a given research question from this peculiar position of digital humanist.

This chapter has also sought to highlight the benefits of the particular methods that I used in my own work, in order to illustrate their effectiveness for answering pieces of a complex research question. For example, the use of network scanning software in repeated passes helped identify the actual changes to network infrastructure across West Africa over a defined period of time. Broad reading practices applied to a large set of public documents helped to identify important texts, and close reading of those particular texts helped to illuminate what a particular telecommunications policy set out to accomplish. Then, moving up a level of abstraction, to higher-order theoretical analysis, articulated a direct relationship between the telecommunications policies and

the network infrastructure changes that we mapped. Most importantly, I arrived at an overall answer to my initial research question precisely by layering all of these methods and ways of thinking. That ability to move between techniques and modes of inquiry, in order to ask original questions, and to answer them, is the great strength of the interdisciplinarity of Digital Humanities.

## REFERENCES

- “Africa Internet Stats Users Telecoms and Population Statistics.” Accessed December 16, 2014. <http://www.internetworldstats.com/africa.htm>.
- “Connecting Africa: ICT Infrastructure Across the Continent.” World Bank, 2010.
- “Homepage | Africa Infrastructure Knowledge Program.” Accessed June 3, 2016. <http://infrastructureafrica.org/>.
- “Index of ftp://ftp.afrinic.net/pub/stats/afrinic/.” Accessed August 11, 2014.
- “Internet Users (per 100 People) | Data | Table.” Accessed March 16, 2014. <http://data.worldbank.org/indicator/IT.NET.USER.P2>.
- “IP Address to Identify Geolocation Information.” IP2Location. Accessed October 15, 2016. <https://www.ip2location.com/>.
- Kaysire, David, and Jiuchang Wei. “ICT Adoption and Usage in Africa: Towards an Efficiency Assessment.” *Information Technology for Development* (September 29, 2015): 1–24. <https://doi.org/10.1080/02681102.2015.1081862>.
- lesbonscomptes. “Recoll Documentation.” Accessed January 21, 2018. <https://www.lesbonscomptes.com/recoll/doc.html>.
- “Projects : West Africa Regional Communications Infrastructure Project—APL-1B | The World Bank.” Accessed January 24, 2016. <http://www.worldbank.org/projects/P122402/west-africa-regional-communications-infrastructure-project-apl-1b?lang=en>.
- Shavitt, Y., E. Shir, and U. Weinsberg. “Near-Deterministic Inference of AS Relationships.” In *10th International Conference on Telecommunications, 2009. Con\_TEL 2009*, 191–198, 2009.
- “The CAIDA AS Relationships Dataset.” Accessed September 20, 2015. <http://www.caida.org/data/as-relationships/>.
- The World Bank. Financing Information and Communication Infrastructure Needs in the Developing World. Public and Private Roles. World Bank Working Paper No. 65. Washington, DC: The World Bank, 2005. [http://wbln0037.worldbank.org/domdoc%5CPRD%5COther%5CPRDDContainer.nsf/All+Documents/85256D2400766CC78525709E005A5B33/\\$File/financingICTreport.pdf?OpenElement](http://wbln0037.worldbank.org/domdoc%5CPRD%5COther%5CPRDDContainer.nsf/All+Documents/85256D2400766CC78525709E005A5B33/$File/financingICTreport.pdf?OpenElement).

“West Africa Regional Communications Infrastructure Program—PID.” Accessed January 23, 2016. <http://www.icafrica.org/en/knowledge-publications/article/west-africa-regional-communications-infrastructure-program-pid-135/>.

World Bank. “Information & Communications Technologies.” Accessed December 14, 2012. <http://web.worldbank.org/WBSITE/EXTERNAL/TOPICS/EXTINFORMATIONANDCOMMUNICATIONANDTECHNOLOGIES/0,,menuPK:282828~pagePK:149018~piPK:149093~theSitePK:282823,00.html>.

“World Development Report 2016: Digital Dividends.” Accessed February 9, 2016. <http://www.worldbank.org/en/publication/wdr2016>.