

Adaptive Electricity Access Planning

by

Stephen James Lee

B.S. Materials Science and Engineering

Second Major in Economics

Johns Hopkins University (2013)

Submitted to the Institute for Data, Systems, and Society
and the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Master of Science in Technology and Policy

and

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018

© Massachusetts Institute of Technology 2018. All rights reserved.

Signature redacted

Author
Institute for Data, Systems, and Society
Department of Electrical Engineering and Computer Science
May 11, 2018

Signature redacted

Certified by
Ignacio J. Pérez-Arriaga
Visiting Professor, Sloan School of Management
Professor, Electrical Engineering, Universidad Pontificia Comillas
Thesis Supervisor

Signature redacted

Certified by
John W. Fisher III
Senior Research Scientist, Computer Science and Artificial Intelligence Laboratory
Thesis Reader

Signature redacted

Accepted by
Munther A. Dahleh
Professor, Electrical Engineering and Computer Science
Director, Institute for Data, Systems, and Society

Signature redacted

Accepted by
Leslie A. Kolodziejski
Professor, Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students



Adaptive Electricity Access Planning

by

Stephen James Lee

Submitted to the Institute for Data, Systems, and Society
and the Department of Electrical Engineering and Computer Science
on May 11, 2018, in partial fulfillment of the
requirements for the degrees of
Master of Science in Technology and Policy
and
Master of Science in Electrical Engineering and Computer Science

Abstract

About 1.1 billion people worldwide lack access to electricity and an additional 1 billion have unreliable access. The social ramifications of this problem are noteworthy because access to electric power has the potential to transform societies. While admirable efforts are underway, there is general consensus that progress is falling far short of what is needed to reach international electricity access goals.

In light of such deficiencies, it is arguable that systems-level experimentation and innovation is required if we are to achieve universal electricity access in the next one to two decades. With the advancement of technology, new opportunities are emerging that can potentially change the game. Machine learning methods and detailed techno-economic models for planning comprise one set of technologies that hold significant promise for accelerating access.

This thesis builds upon recent work towards the development of more intelligent decision support systems for electrification planning. Progress towards automated and scalable software systems for the extraction of building footprints from satellite imagery are presented. In addition, a novel model for probabilistic data fusion and other machine learning methods are compared for electrification status estimation. Inference tools such as these allow for the cost-effective provision of granular data required by techno-economic models.

We also acknowledge that the technologies we detail should not be developed in a vacuum. Given that electrification is a complex endeavor involving numerous social and technical factors, careful consideration must be given to human, policy, and regulatory concerns during the planning process. We notice how uncertainty abounds in these activities and propose “adaptive electricity access planning” as a new model-assisted framework for the explicit consideration of uncertainty in large-scale planning. This work aspires to provide valuable perspective on the importance of uncertainty in planning as these endeavors continue to evolve.

Thesis Supervisor: Ignacio J. Pérez-Arriaga

Title: Visiting Professor, Sloan School of Management

Professor, Electrical Engineering, Universidad Pontificia Comillas

Thesis Reader: John W. Fisher III

Title: Senior Research Scientist, Computer Science and Artificial Intelligence Laboratory

Acknowledgments

To me, this thesis represents the culmination of an incredibly rewarding nearly three year journey. My time with the Universal Energy Access Lab, TPP, and the SLI Group was spent meeting amazing people, forming friendships, studying diverse subjects, exploring four continents, learning, and living. For supporting the development of my research and much more, I am sincerely grateful to:

- Ignacio Perez-Arriaga, for his guidance, humor, and complimentary Spanish lessons. His expertise regarding power systems, economics, and regulation were invaluable to my understanding of how energy powers the world, and his passion for positive social change is contagious. ¡Muchas gracias para las tapas ricas, todas las risas, y su liderazgo en Patna (la nuit)!
- John Fisher, for welcoming me into the Sensing, Learning, and Inference Group and introducing me to the world of generative models. Thanks for the tips on building models “the right way!”
- Claudio Vergara, for introducing me to power systems planning with REM, showing me that MATLAB isn’t actually all that bad, and for the climbing lessons between research meetings.
- Chris Dean, for teaching me about information planning and inference on probabilistic graphical models, and for never running out of quips and math puns.
- George Chen and Bolei Zhou, for introducing me to deep learning, helping me to work with Turkers, and for their overall mentorship.
- Rob Stoner, for his perspectives on entrepreneurship and economic development, for his research guidance, and for sharing stories of rural Canada.
- Larry McCray and Ken Oye, for the vivid science policy discussions and introducing me to planned adaptation: by far my favorite policy tool.
- The rest of the UEA Lab for sharing in journeys around the world and for contributing to something I’ve found special. Cailinn, Turner, Vivian, Roxanne, Olamide, Matt, Shweta, Patricia, Simone, and Andres: thanks for helping to keep me alive while on the road! Thanks also to Pablo, Yael, Luca, Rafael, Reja, Pedro, Fernando, Carlos, Gregoire, and Marc.
- The rest of the SLI Group - Randi, David, Sue, Jason, Vadim, Rujian, Genevieve, and Julian - for showing me how to be a better Bayesian and prepare for the lentil revolution.
- Our group’s friendly image annotators: Meghan, Theresa, Mary, Frances, and Tomas; our partners at GIZ, Tata Power DDL, Tata Trusts, REA Uganda, NBPDC, EDCL, Development Seed, and the Shell Foundation; and the kind people from the Bihari villages who shared with us their ways of life.

- MITEI, the Tata Center, 6.867 squad, Frank, Barb, and the TPP community.

and lastly, my friends and family. I especially want to thank my parents for their encouragement and support all these years.

Contents

Contents	7
List of Figures	13
List of Tables	15
1 Introduction	17
2 Universal Access to Electricity	21
2.1 Why Care About Electricity Access?	22
2.2 Socioeconomic Benefits of Electrification	22
2.2.1 Poverty and Inequality	23
2.2.2 Climate Change and Electricity Access	25
2.3 Technologies for Electrification	26
2.3.1 Grid Electrification	26
2.3.2 Mini-grids	26
2.3.3 Isolated Units	27
2.4 Defining Electricity Access	27
2.5 The Economics of Electricity Access	28
2.5.1 Demand Density and Supply Costs	30
2.5.2 Consumer Demand and Diminishing Marginal Utility	30
2.5.3 Subsidization and the Viability Gap	31
2.6 Political Factors and Consumer Preferences	32
2.6.1 Politicization	32
2.6.2 Corruption	33
2.7 Regulation and Business Models	33
2.8 Conclusions	34
3 Building Footprint Extraction and Localization	35
3.1 Background	36
3.1.1 Feedforward Neural Networks	36
3.1.2 Convolutional Neural Networks	38
3.1.3 Traditional Methods for Automatic Satellite and Aerial Image Labeling	40

3.1.4	Convolutional Neural Networks for Automatic Satellite and Aerial Image Labeling	41
3.1.5	Satellite imagery	42
3.1.6	Convolutional Neural Networks for Semantic Segmentation	42
3.2	Procuring Training Data	44
3.2.1	Using Microsoft Paint and Google Earth	44
3.2.2	Amazon Mechanical Turk	45
	TurkCleaner Tests on Amazon Mechanical Turk	45
	DrawMe Tests on Amazon Mechanical Turk	46
	DrawMe and Back End System Tests on Amazon Mechanical Turk	48
3.2.3	Using Web Tools, Google’s Realtime API, freelancer.com, and Expert Annotators	49
3.3	Building Extraction and Image Quality Studies	50
3.3.1	Data Procurement and Formatting	50
3.3.2	Error Metrics	51
3.3.3	Assessment of FCN Semantic Segmentation Quality	53
3.3.4	Image Quality	54
3.3.5	Data Augmentation	55
3.4	Load Localization and Characterization	56
3.5	Large Scale Building and Load Localization	59
3.6	Conclusions	60
4	Electrification Status Estimation	61
4.1	Background	62
4.1.1	Electrification Status Estimation	62
4.1.2	Nighttime Lights	64
4.1.3	Distance to Transformers and Electrification	64
4.1.4	Hierarchical Spatial Models and Hierarchical Models	65
4.2	Features Related to Electrification	65
4.2.1	Global Inputs	65
	Nighttime Lights Data	65
4.2.2	Inputs for the South Service Territory, Uganda	66
	Building Location Data	66
	Transformer Locations and Distances, 2016	67
	Sparse Survey Data, 2016	68
4.2.3	Inputs for Uganda at the Country-Level	69
	Aggregate Census Data, 2014	69
	Transformer Locations, 2011	70
	Sparse Survey Data, 2012	70
4.2.4	Inputs for for Kayonza, Rwanda	71
	Building Location Data	71
	Transformer Locations, 2017	72
	Low Voltage Network	72
	Aggregate Census Data and Derivatives for Electrified Buildings	72

4.3	Electrification Status Models and Error Rates	73
4.4	A Logistic Regression Model	74
4.4.1	Background	74
	The Bernoulli Distribution	74
	Logistic regression	74
4.4.2	Uganda Case Study	75
4.4.3	South Service Territory, Uganda Case Study	77
4.5	A Gaussian Process Model	78
4.5.1	Background	79
4.5.2	Methods	81
4.5.3	Uganda Case Study	82
4.5.4	South Service Territory, Uganda Case Study	84
4.5.5	Discussion	86
4.6	The Hierarchical Beta Model	86
4.6.1	Background	87
	Bayes' Rule	87
	The Binomial Distribution	88
	The Beta Distribution	88
	The Gamma Distribution	89
	Probabilistic Graphical Models	90
	Directed Graphical Models	90
	Inference	91
	Sampling and Markov Chain Monte Carlo Methods	91
	Metropolis Hastings Algorithm	92
	Gibbs Sampling	93
	Hierarchical Bayes	94
	Empirical Bayes	94
4.6.2	Methods	94
	Enforcing Multiscale Smoothness	95
	Enforcing Spatial Smoothness	97
	Function Approximation with Empirical Bayes	99
	Inference	102
4.6.3	Case Studies	104
4.6.4	Discussion	106
4.7	Conclusions	106
5	Planning and Techno-Economic Models	109
5.1	Electrification Planning	109
5.2	Techno-Economic Models	111
5.2.1	Cost-Optimal Planning Considerations	112
	Grid Electrification	112
	Off-grid Mini-grids and Isolated Units	112
	The Feasible Region	113
5.3	The Reference Electrification Model	114
5.3.1	Off-Grid Advantage and Grid Extension Advantage	116

5.4	Conclusion	117
6	Electricity Access Planning Under Uncertainty	119
6.1	Uncertainty	121
6.1.1	Types and Sources of Uncertainty	122
	Sources of Uncertainty Relevant to Modeling	122
	Knowledge and Metaknowledge	123
	Quantifiable and Unquantifiable Uncertainty	123
6.1.2	Sources of Uncertainty in Electricity Access Planning	124
6.2	Adaptive Electricity Access Planning	125
6.2.1	Electrification Prioritization under Uncertainty	127
6.3	Dealing with Quantifiable Uncertainty	131
6.3.1	Information Planning and Decision-Making under Uncertainty	131
6.3.2	Flexible Design and Real Options	132
6.4	Dealing with Unquantifiable Uncertainty	132
6.5	Platforms for Coordination	134
6.6	Conclusions	134
7	Conclusions	137
	Bibliography	139

Abbreviations

ConvNet Convolutional neural network

CRF Conditional random field

DAG Directed acyclic graph

DMSP-OLS Defense Meteorological Satellite Program - Operational Linescan System

EDCL Energy Development Corporation Limited, Rwanda

FCN Fully connected network

FITC Fully independent training conditional approximation

GIS Geographic information system

GIZ Deutsche Gesellschaft für Internationale Zusammenarbeit GmbH

GP Gaussian process

GPC Gaussian process classification

HIT Human Intelligence Task

IEA International Energy Agency

IEG Independent Evaluation Group, World Bank

k-NN k-Nearest neighbors algorithm

LA Laplace approximation

MCMC Markov chain Monte Carlo

MH Metropolis Hastings

MTF Multi-tier framework

MTurk Amazon Mechanical Turk

NOAA National Oceanic and Atmospheric Administration

OnSSET Open Source Spatial Electrification Toolkit

OSM OpenStreetMap

pdf Probability density function

PDIA Problem-Driven Iterative Adaptation

pmf Probability mass function
PV Photovoltaic
REA Uganda Rural Electrification Agency of Uganda
REM Reference Electrification Model
SDG 7 Sustainable Development Goal #7
SEforALL The Sustainable Energy for All Initiative
SGD Stochastic gradient descent
SVM Support vector machine
UBOS Uganda Bureau of Statistics
Uganda SST South Service Territory, Uganda
UN United Nations
VoI Value of information
WEO World Energy Outlook

List of Figures

2-1	Probable pathways for benefits from electricity access	22
2-2	World Bank multi-tier ratings	28
2-3	SEforAll multi-tier matrices for household access to electricity	29
2-4	Consumer Demand	31
3-1	Simple feedforward neural network	37
3-2	Sparse feedforward neural network with skip layer connections	38
3-3	Layers in convolutional neural networks	39
3-4	Model architectures for fully convolutional networks for semantic segmentation	43
3-5	The TurkCleaner interface	45
3-6	Sample MTurk annotations	47
3-7	Database schema for MTurk back end tool	49
3-8	Modified DrawMe interface	50
3-9	Ground truth for building extraction	51
3-10	FCN inferences	53
3-11	Learning curves for FCN runs	55
3-12	Load localization approaches	57
3-13	Large scale building identification for Vaishali, Bihar, India (>2000 km ²)	59
4-1	2013 DMSP-OLS annual composite imagery	66
4-2	2013 DMSP-OLS annual composite imagery for the Uganda SST	67
4-3	Uganda SST building locations	68
4-4	Uganda SST 2016 transformers and sub-country borders	69
4-5	Uganda SST 2016 electrification survey	70
4-6	Uganda 2012 country-level survey, 2013 nighttime lights, and 2011 transformer data	71
4-7	Low voltage grid, transformer locations, and buildings for Kayonza, Rwanda	72
4-8	Logistic regression separating lines for the Uganda case study	76
4-9	Uganda SST survey locations, nightlights, and transformers	78
4-10	Logistic regression separating line for the Uganda SST case study	79
4-11	GP contour plots for Uganda	83
4-12	GP results corresponding to the 2012 ERT survey, zoomed into the capitol city, Kampala.	84

4-13	GP results corresponding to the 2016 REA survey for the Uganda SST	85
4-14	Beta and gamma distribution examples	89
4-15	General representation of the hierarchical beta model	95
4-16	Plate notation representation of the hierarchical beta model	96
4-17	Extended graphical representation of the hierarchical beta model showing spatial smoothness	98
4-18	Smoothness demonstrations for the hierarchical beta model	99
4-19	Empirical Bayes in the hierarchical beta model	100
4-20	Conditional independencies enabling efficient inference	103
4-21	Uganda SST case study for the hierarchical beta model	104
4-22	Kayonza case study for the hierarchical beta model	105
5-1	The SEAP framework	111
5-2	A framework for thinking about techno-economic models for electrification planning	113
5-3	Web visualizations for REM output	115
5-4	REM Off-Grid Advantage heatmap	116
6-1	Rwandan electrification prioritization from 2009	120
6-2	Flow diagram for an adaptive electricity access plan	126
6-3	Uncertainty approximations for electrification projects	128
6-4	Uncertainty thresholds for investment consideration	129
6-5	Investment impact prioritization	130

List of Tables

3.1	Semantic segmentation metrics	54
4.1	Methods for electrification status estimation	62
4.2	Logistic regression error metrics for the 2012 Uganda case study . . .	77
4.3	Logistic regression error metrics for the 2016 Uganda SST case study	79
4.4	GP compared to other classifiers for the 2012 Uganda case study . . .	82
4.5	GP compared to all-negative classifier for the 2016 Uganda SST case study	85
6.1	Johari window variant	123

List of Algorithms

1	Error backpropagation algorithm	38
2	Load localization with exclusion zones	58
3	Metropolis Hastings algorithm	93
4	Gibbs sampler algorithm	93

Chapter 1

Introduction

About 1.1 billion people worldwide lack access to electricity and an additional 1 billion have unreliable access (United Nations Foundation, 2017). The social ramifications of this problem are massive because access to electric power has the potential to change lives. The provision of even just lighting, one of the most basic electric services, has been shown to increase evening study and productive work hours, improve educational outcomes and access to opportunity, reduce household air pollution from kerosene lamp substitution, enhance public security, and improve quality of life (Banerjee et al., 2014). Additionally, modern energy access enables important services such as household cooling, communication, entertainment, water pumping, and refrigeration (World Energy Outlook, 2017).

Though much of the world experienced tremendous progress in energy access over the last century, the challenge of energy poverty is still immense. The vast majority of those without modern energy access today reside in South Asia and sub-Saharan Africa and are also among the world's poorest people. In the energy poverty literature, today's populations without access are considered the "last mile." For myriad reasons including economic constraints and institutional challenges, providing energy access in these regions has proven more difficult than elsewhere (Halff et al., 2014). While the UN Sustainable Development Goal #7 (SDG 7) targets "universal access to affordable, reliable and modern energy services" by the year 2030, there is general consensus that progress to-date is falling short of what is needed (United Nations, 2017). The IEA's 2017 WEO report projects that, under the status quo "New Policies Scenario," 675 MM will still lack electricity access by 2030 and this number will grow to 711 MM by 2040 due to population growth in sub-Saharan Africa (International Energy Agency, 2017). It is also projected that a significant number of those with access will have unreliable supplies (Independent Evaluation Group, The World Bank Group, 2015).

Planning for universal energy access is a complex endeavor, exemplifying coordination issues between multiple stakeholders, technological change, business model innovation, local environmental factors, global climate and health concerns, perceptual and cultural differences, and financing challenges. These considerations give effective energy access planning the potential to be highly diverse over time and space; countries today can take different routes toward universal access from those in the past, and can equivalently take different routes from one another going forward.

Energy planning models and toolkits have been developed to cut through some of these complexities and support concrete and effective planning decisions. At a high level, these models inform technology choice through power system design, support regulatory and policy decisions, and promote stakeholder coordination. Nevertheless, model detail and accuracy is fundamentally constrained by input data availability and granularity.

This thesis covers ways that modern machine learning methods can assist in improving pertinent input data and thus enable detailed electrification planning models. It also provides new perspectives about how such models can fit into adaptive frameworks for iterative and self-correcting plans.

Chapter 2 provides an overview of universal energy access with discussions of economic, political, and regulatory factors relevant to electrification planning. Subsequently, capabilities for extracting building footprints from satellite imagery and associated studies are presented in chapter 3, and new approaches for estimating building electrification status are presented in chapter 4. With the cost-effective provision of more granular data on building locations and electrification status, the feasibility of performing accurate household-level planning improves substantially. We describe electrification planning and contemporary techno-economic models for cost-optimal planning in chapter 5. In chapter 6, we expound on this discussion by proposing a framework for combining characterizations of uncertainty to inform decision-making for information and infrastructure investments. Such measures allow for highly iterative planning and can be related to various research areas for decision-making under uncertainty. Through these discussions, we present the following contributions:

- Describing methods for building footprint extraction from satellite imagery and presenting their application to large-scale energy access planning.
- Accounting for challenges encountered designing software systems for the procurement of building footprint extraction training data and describing system designs that have been found to work well in practice.
- Presenting electrification status estimation as a new application of machine learning techniques including Gaussian processes and Bayesian networks.
- Describing the hierarchical beta model: a novel approach for efficient multi-modal and multiscale data fusion for spatial processes.
- Providing an overview of types of uncertainty that are relevant to electrification planning, and differentiating between quantifiable and unquantifiable sources of uncertainty.
- Sharing perspectives on how techno-economic models such as the Reference Electrification Model can inform project prioritization efforts using quantified measures of uncertainty and estimated costs for a set of supply technologies.
- Proposing “adaptive electricity access planning” as a new model-assisted framework for the explicit consideration of uncertainty in decision-making processes

for electrification. Adaptive electricity access plans emphasize closed-loop and iterative designs for concurrent information and infrastructure planning. Pathways for value creation are also outlined by discriminating between quantifiable and unquantifiable forms of uncertainty.

- Elucidating how modeling and information planning approaches from the machine learning literature have the potential to help planners quantify uncertainties, maximize the utility of scarce geospatial data, and plan for the efficient procurement of such data under resource constraints.
- Describing how notions of flexible design from the engineering systems literature can facilitate the nuanced but likely impactful consideration of uncertain futures for infrastructure design.
- Highlighting that unquantifiable uncertainties persist regardless of modeling endeavors and underscoring the importance of adaptability in planning and decision-making procedures. Relating electricity access planning to policy tools such as planned adaptation, adaptive management, and Problem-Driven Iterative Adaptation from the law and political science literatures.
- Presenting opportunities for the beneficial integration of adaptive electricity access plans and platforms for collaboration in terms of reducing information asymmetries between stakeholders.

Each of the first five chapters is largely self-contained, though they are altogether related and provide requisite background for the proposal of “adaptive electricity access plans” as presented in Chapter 6. While the presentation of adaptive electricity access planning in this thesis comprises a wide-ranging framework for information and infrastructure planning under uncertainty, the proposal presented in this thesis represents only a first iteration towards a workable design. Further vetting of lower-level details are required and the full costs of pursuing such an approach may only be apparent if the framework is actually implemented. Nevertheless, it is hoped that the proposal and methods described in this thesis provide a valuable perspective on electrification planning under uncertainty. Uncertainty abounds in electrification planning, which is an activity that can be transformative for whole societies. Explicitly incorporating uncertainty into increasingly model-driven planning frameworks holds potential worthy of consideration.

Chapter 2

Universal Access to Electricity

Universal energy access has received increasing levels of attention due to consensus about its importance as a key enabler for economic development. Exemplifying this, the United Nations (UN) designated 2012 as the Year of Sustainable Energy for All, and in 2014, it began the Decade of Sustainable Energy for All. Also in 2014, the UN Secretary-General announced the Sustainable Energy for All initiative (SEforALL), which endeavors to support Sustainable Development Goal #7 (SDG 7), calling for the attainment of “affordable, reliable, sustainable and modern energy for all by 2030” (Birol, 2014; Sustainable Energy for All, 2017).

Nevertheless, many recognize that we are not progressing rapidly enough to meet SDG 7’s universal access targets. In 2014, the International Energy Agency (IEA) produced a World Energy Outlook (WEO) Special Report entitled “Africa Energy Outlook: A Focus on Energy Prospects in Sub-Saharan Africa.” The report provides projections for progress in electrification for Sub-Saharan Africa out to 2040, and predicts that although nearly a billion Africans will gain access to electricity in this timespan, more than a half billion will remain disconnected due to population growth (International Energy Agency, 2014). The Independent Evaluation Group (IEG) within the World Bank shares this pessimistic outlook in their report, “World Bank Group Support to Electricity Access, FY2000-2014.” They find that by extrapolating electrification trends from 2000-2010 and including expected population growth, the non-electrified population will actually increase from 1.1 billion today to 1.2 billion by 2030 worldwide. Additionally, they predict that 1 billion will still lack reliable supplies as of that date (World Bank Independent Evaluation Group, 2015). The IEA’s 2017 WEO report is slightly more optimistic and projects that, under the status quo “New Policies Scenario,” 675 MM will still lack electricity access by 2030 and this number will grow to 711 MM by 2040 due to population growth in sub-Saharan Africa (International Energy Agency, 2017). In any case, progress has been and is projected to be significantly behind where it needs to be.

Because status quo interventions will likely not extend electricity access fast enough to meet international goals, technological change, regulatory innovation, and the development of new business models appear critical to improving upon historical rates of progress. This chapter presents reasons for why universal energy and electricity access is important, discusses technologies for electrification, examines definitions

of electricity access, and provides a brief overview of electrification economics. The purpose of this chapter is to outline the importance of developing methods to support electricity access planning and to provide background for subsequent discussions about adaptive management and planning.

2.1 Why Care About Electricity Access?

Electricity access and energy access in general have been described as important social goals due to the myriad socioeconomic benefits they promote and their associated ramifications towards the reduction of poverty and inequality. Intelligent planning may also enable the joint consideration of related social goals, including long term power sector decarbonization.

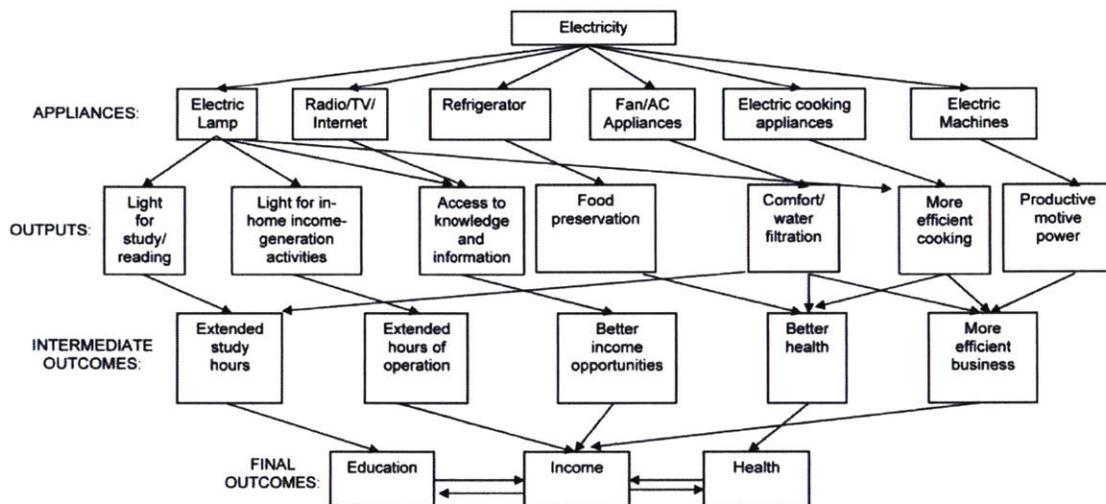


Figure 2-1: Probable pathways for benefits from electricity access. Figure source: (Khandker et al., 2013)

2.2 Socioeconomic Benefits of Electrification

Electrification is commonly understood to confer significant socioeconomic benefits in both direct and indirect ways. The provision of electric lighting has been shown to prolong evening study and productive work hours, improve educational outcomes, and enhance access to opportunity. Furthermore, substituting polluting kerosene lamps with electric lighting reduces risk of respiratory disease among women and children (Banerjee et al., 2014). Outside of lighting, electricity access can improve quality of life in numerous ways. Khandker et al. describe probable pathways for how electricity can impact education, income, and health through the employment of value-adding appliances. The diagram in Fig. 2-1 shows how appliances like radios, television, Internet-connected devices, refrigerators, fans and air conditioning, electric cooking

appliances, and electric machines eventually lead to benefits including extended study and productive work hours, improved income opportunities, better health, and more efficient business (Khandker et al., 2013).

The literature suggests that the effects of some of these pathways are profound. Improved educational outcomes are critical to the long-term development of human capital within communities. In addition, the ability to access information through radios, television, and Internet-connected devices gives family members, especially women, access to knowledge which can elevate their household decision-making power (Banerjee et al., 2014). Furthermore, electric irrigation pumps can also significantly improve farm productivity (Khandker et al., 2013). Ultimately, electricity access may enable human empowerment and improved quality of life (Practical Action, 2016).

Despite clear anecdotal and logical evidence for the socioeconomic benefits of electrification, empirical studies regarding these benefits have been more mixed. Khandker et al. use panel-estimation studies on energy access in Vietnam using 2002 and 2005 surveys and find that household electrification has increased incomes by 28%, expenditures by 23%, school attendance for girls by 9%, and school attendance for boys by 6.3% (Khandker et al., 2013). In contrast, Burlig et al. use a regression discontinuity analysis to estimate the effects of rural electrification on economic development in India. Nighttime lights, administrative boundary data, and census information are combined to compare otherwise identical non-electrified and electrified villages. Burlig et al. find that while electricity use increases in electrified villages, no substantial benefits accrue with regards to labor markets, asset ownership, housing quality, or education. The only perceptible shift was in men moving from agricultural to non-agricultural modes of employment (Burlig and Preonas, 2016). Corroborating this finding, Aklin et al. conduct a randomized field experiment on 1,281 Indian rural households and find that an increase in off-grid electrification rates from 29% to 36% within one year resulted in no systematic changes in savings, spending, business creation, or time spent working and studying (Aklin et al., 2017).

A more holistic view may be required to characterize the benefits of electrification. Halff et al. share the perspective that energy access is necessary but not sufficient to guarantee economic development (Halff et al., 2014). Electricity access is just one enabling factor; other factors include means of transportation to import and export goods, sanitation, education, and some potential for productive uses of electricity. Economic development is a complex endeavor and is dependent upon interrelated contributing factors. While energy access is very important, it is not in itself a panacea.

2.2.1 Poverty and Inequality

Energy poverty is fundamentally linked with the economic concepts of income poverty and inequality. Those without access have intrinsic restrictions to their earnings opportunities and consumption possibilities. A few concepts are often cited for the consideration of equality as a social goal, and it may be argued that they relate to energy poverty the same way as they relate to income or wealth.

The first concept is predicated on the notion of fairness. In the case of electri-

fication, this can be framed as a question: why should some members of a country benefit from the provision of public infrastructure and electricity services while others shouldn't? The underlying argument is that citizens should be entitled to a just share of their society's wealth just for being members of it (Wolff, 2009). While some perceive that the notion of fairness should only concern fairness in opportunity and not wealth, it can be argued that the provision of electricity access affects both opportunity and wealth.

A related concept concerns John Rawls', "A Theory of Justice." Rawls provides a thought experiment based on a person's "original position." He argues that a person behind a "veil of ignorance" who does not know his or her income or wealth would prefer to live a society that provided a minimal level of well-being to all of its members (Rawls, 1971). Rawls' views can be extrapolated to the idea of a "social safety net," where governments provide minimum levels of income and necessary services to its citizens. Amartya Sen expands on this view in his work on inequality and emphasizes the importance of considering equality pertaining to "functionings." Elementary functionings pertain to good health, nourishment, and shelter, while more complex social functionings include community involvement and self-respect (Sen, 1992; Wolff, 2009). Sen also distinguishes between "resources," like income and wealth, and "capabilities," for which energy access is an example. Capabilities pertain to what individuals are able to do and are distinct from resources. For instance, a person with significant wealth and the latest Tesla Model S sedan may be less "capable" than a person with a camel, if she is without access to the road and fueling infrastructure her Tesla requires. Sen argues for a "capability approach" over a "resource-centered" one, as capabilities are more linked with "substantive freedom" and may better characterize poverty than measures of income alone (Halff et al., 2014).

Another reason for equality as a social goal is based on philosopher Jeremy Bentham's ethical theory of utilitarianism. Bentham makes the normative judgment that a society's goal should be to maximize the total utility of its citizens. With the assumption of diminishing marginal utility and the assumption that all members of society have the same utility functions with regards to income, it is straightforward to show that total utility is maximized when incomes are evenly distributed (Bentham, 1823; Wolff, 2009). While this argument is likely imperfect due to flaws in the aforementioned assumptions and the potential disincentives to work and invest associated with redistributions of income, it is arguable that the utility benefits of universal energy access are likely very positive on net.

Finally, a simple argument against poverty and inequality concerns their associated negative externalities. Poverty and inequality may result in an increase in the number of homeless people, beggars, and criminals in society (Wolff, 2009). In this way, redistribution and the subsidization of electric power may be in everyone's direct self-interest. This is especially true over longer time horizons, when the full economic benefits of expenditures on electrification have had time to manifest.

2.2.2 Climate Change and Electricity Access

The relationships between energy access and climate change are subject to ongoing debate. The IEA estimated that achieving universal energy access (including the provision of clean cooking facilities) by 2030 would cost \$49 billion per year, starting in 2011. While this number is five times greater than prior investment figures, it still accounts for only 3% of expected worldwide energy infrastructure investments over this period. The IEA further estimates that achieving universal access will increase electricity generation by only 2.5%, fossil fuel demand by 1%, and carbon dioxide emissions by 0.6% (Biro, 2014).

Despite these seemingly reassuring figures, uncertainty around the real costs of electrification abound. The IEA's estimates rest upon the assumptions that newly connected people will have low levels of energy consumption and that electrification will occur via a high share of renewables. The challenge with the former assumption is that mechanisms are still unclear regarding how energy consumption grows once it is delivered. In addition, since economic development is ultimately desired, simply calculating consumption characteristics until basic access is achieved is inadequate. Higher levels of consumption are required over time to enable the full benefits of development. The latter assumption is problematic as well since technological choices are difficult to forecast. Implementation types depend on local preferences, plans, policies, market conditions, and technological developments. Complicated trade-offs persist between climate goals, economic development, and electrification modes and are active areas of continued research.

One nuanced consideration regarding electrification mode and decarbonization pertains to the lifecycle developments of off-grid solutions. As the next section details, electrification through grid extension is often more polluting than that using mini-grids due to its likely greater reliance on fossil fuel-powered generation sources. In addition, since mini-grids commonly serve smaller loads and are deployed in remote areas, line extension costs often become prohibitive making renewables the most economical generation option. Specifications for mini-grids vary significantly and there is an important distinction to be made between low-quality grid-incompatible and higher-quality grid-compatible systems. Since mini-grids often represent transient modes of electrification for rural villages that will ultimately connect to the main grid, grid-incompatible mini-grids are fully discarded when the grid arrives to these service areas. This usually results in the abandonment of renewable generation assets as well as yielding significant investor risk. In this scenario, villages that were once powered by clean energy resources will instead take on the grid's relatively more polluting generation mix. If up front investments are made so that mini-grids are instead designed to provide the option for grid-connection, mini-grid infrastructure may be migrated and duplicative investments can be avoided when the main grid arrives. This results in investor risk reduction as well as the preservation of renewable generation assets. Considerations like these motivate elevated conversation about electrification pathways and may require detailed techno-economic modeling endeavors for high level decision-making.

2.3 Technologies for Electrification

As alluded to in the previous section, numerous competing technologies exist for household electrification. Grid extension, mini-grids, and isolated systems may reflect optimal electrification modes in different contexts; their relative attractiveness depends on factors including customer preferences, affordability, energy resource availability, cost, expectations for the future, and network effects. In some cases, it might even make sense for complementary modes to be pursued simultaneously. The IEA’s 2017 WEO report projects that under the status quo “New Policies Scenario” the world’s population without electricity will fall to 675 MM by 2030. While progress in urban electrification will be dominated by new grid connections, only 30% of progress in rural electrification will come from the grid. 37% of this number will be supplied by mini-grids and isolated off-grid solutions will account for the remaining 33% (International Energy Agency, 2017). Careful consideration regarding technological choice is necessary to ensure expenditures on electrification yield worthwhile societal benefit when constrained by resource scarcity.

2.3.1 Grid Electrification

Connection to a main electrical grid has traditionally been the status quo mode of electrification in both developed and developing world. Large generation facilities benefit from economies of scale and a network of transmission and distribution lines deliver power to consumers. The main grid is usually part of large national or regional interconnected systems (Pérez-Arriaga, 2017). In order for non-electrified buildings to connect to the grid, it must be close to a power line. Because extending the grid over long distances requires new infrastructure to be built, it is generally much more expensive to connect a building that is far away from existing lines than one that is nearby (Lee et al., 2016).

2.3.2 Mini-grids

Mini-grids and microgrids are smaller power systems that connect buildings within a localized region for energy generation, distribution, and oftentimes storage. Though the names ‘mini-grid’ and ‘microgrid’ are sometimes used to distinguish between system sizes (e.g., mini-grids may be defined as systems with 100 kW to 10 MW of installed capacity and microgrids with less than 100 kW, etc.), they are also often used without distinction (Sovacool, 2014; Pérez-Arriaga, 2017). For simplicity, we use the term ‘mini-grid’ generically. Mini-grids are usually managed locally and can be powered by solar PV, biomass combustion, fossil fuel-powered generation (e.g., diesel), micro-hydro, and wind turbines. More advanced mini-grids may be designed to operate alongside the centralized grid; they can interconnect when it is advantageous to buy and sell power to the centralized grid and disconnect to maintain power quality. Mini-grids are generally economically advantageous over central grid extension in remote regions where it is prohibitively expensive to extend the grid; however, they may also represent sensible investments in more densely populated areas if the

central grid is highly unreliable. This latter case is often related to distribution company insolvency. When insolvent, a distribution company may be financially unable to invest in grid extensions, leaving alternatively financed mini-grid electrification to be the next-best solution.

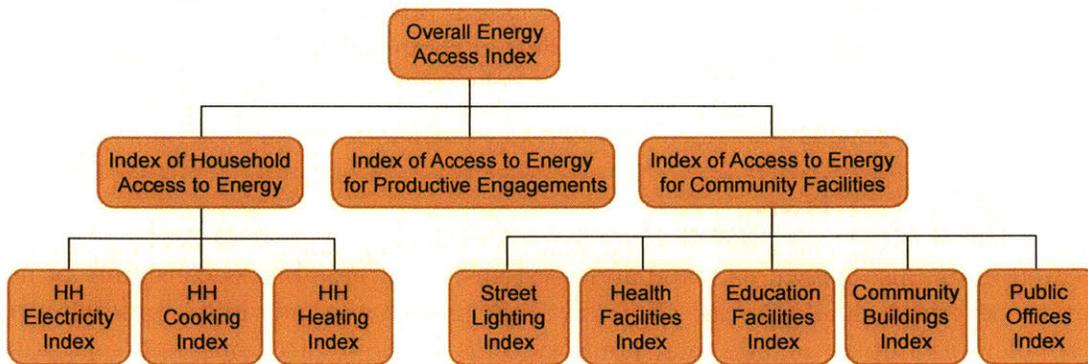
2.3.3 Isolated Units

Isolated units, also known as standalone systems, are a broad class consisting of numerous off-grid technologies for electricity provision at the individual building-level or smaller. Examples include solar home systems, micro-hydro dams, and solar kits. The type of isolated unit that should be employed depends on available fuel supplies, proximity to neighboring consumers, and affordability. Solar kits, for instance, are relatively inexpensive but provide rudimentary forms of electricity access relative to other technologies. Micro-hydro dams and solar home systems have higher up-front costs and are only sensible in areas with running water and sufficient sunlight, respectively (Sovacool, 2014; Pérez-Arriaga, 2017).

2.4 Defining Electricity Access

Defining electricity access is a more complicated endeavor than it may initially seem. Binary metrics for electrification status have been used traditionally. Given a binary metric, electricity is considered present if a household is connected to the grid and it is not otherwise. Binary metrics have several advantages including ease of interpretability and ability to be aggregated in a straightforward way. For instance, it is easy to understand a statement like, “as of 2014, 19.8% of the Rwandan population had access to electricity.” In an ideal system, this metric would be very illuminating; however, in the presence of low power reliability, affordability constraints to consumption, and illegal connections, binary metrics hide important information regarding the value that people derive from electric power (Bhatia and Angelou, 2015). For instance, if there is very low reliability, a household may not derive any benefits from grid connection for most of the day. If outages are unpredictable or unannounced, the ability to use electricity towards productive ends decreases as well. Affordability constraints pose similar difficulties for the derivation of value from grid connection; if a household cannot afford basic services, they may as well not be connected. Finally, illegal connections can be misleading. While they may be considered connections, they deteriorate the financial health of the utility and make it more difficult for utilities to grow and sustainably extend the grid. Illegal connections may also represent public safety hazards.

As a result of deficiencies associated with binary metrics for electrification status, SEforAll developed what they call the Multi-tier Framework (MTF) for household electricity access. They go further to define the “Overall Energy Access Index” as depicted in Fig. 2-2. The overall index is composed of three “locales of energy access,” which are represented by the “Index of Household Access to Energy,” “Index of Access to Energy for Productive Engagements,” and “Index of Access to Energy for



Note | HH = household

Figure 2-2: World Bank multi-tier ratings. The rating for a given household is equal to the lowest tier-rating across all attributes. Figure source: (Bhatia and Angelou, 2015)

Community Facilities.” As can be seen in the tree diagram, each of the locales are further divided into sub-locales. The “Household Electricity Index” is an example of a sub-locale within the “Index of Household Access to Energy.” Within the “Household Electricity Index,” SEforAll provides multi-tier matrices for measuring access to household electricity supply, household electricity services, and household electricity consumption, as shown in Fig. 2-3a, 2-3b, and 2-3c. One or many attributes are given by the matrices, and the multi-tier rating for a given household ranges from 0 to 5 and is equal to the lowest tier-rating across all attributes. Tier ratings can then be aggregated and scaled to provide higher-level index values. Household electricity supply is particularly multifaceted; it takes account of attributes including peak capacity, availability, reliability, quality, affordability, legality, and safety (Bhatia and Angelou, 2015). The spirit behind SEforAll’s MTF is to provide more granular information about the state of electricity access in developing countries relative to binary metrics. Detailed descriptions of electrification status such as those provided by the MTF enable countries to make more truthful assessments of their progress and help them to infer what the most cost effective interventions may be.

2.5 The Economics of Electricity Access

The economics of electricity access are critical to thinking about planning. In order for meaningful electrification to take place, infrastructure must be built and power needs to flow to end consumers. For grid extension, poles must be erected and lines installed and reinforced. In addition, upstream supply must account for increased demand from the service of new customers, typically by means of new generation capacity.

		TIER 0	TIER 1	TIER 2	TIER 3	TIER 4	TIER 5
ATTRIBUTES	1. Peak Capacity	Power capacity ratings ²⁸ (in W or daily Wh)	Min 3 W	Min 50 W	Min 200 W	Min 800 W	Min 2 kW
			Min 12 Wh	Min 200 Wh	Min 1.0 kWh	Min 3.4 kWh	Min 8.2 kWh
		OR Services	Lighting of 1,000 lmhr/day	Electrical lighting, air circulation, television, and phone charging are possible			
	2. Availability (Duration)	Hours per day	Min 4 hrs	Min 4 hrs	Min 8 hrs	Min 16 hrs	Min 23 hrs
		Hours per evening	Min 1 hr	Min 2 hrs	Min 3 hrs	Min 4 hrs	Min 4 hrs
	3. Reliability					Max 14 disruptions per week	Max 3 disruptions per week of total duration <2 hrs
	4. Quality					Voltage problems do not affect the use of desired appliances	
	5. Affordability				Cost of a standard consumption package of 365 kWh/year < 5% of household income		
6. Legality					Bill is paid to the utility, pre-paid card seller, or authorized representative		
7. Health & Safety					Absence of past accidents and perception of high risk in the future		

(a)

	TIER 0	TIER 1	TIER 2	TIER 3	TIER 4	TIER 5
Tier criteria		Task lighting AND Phone charging	General lighting AND Phone Charging AND Television AND Fan (if needed)	Tier 2 AND Any medium-power appliances	Tier 3 AND Any high-power appliances	Tier 2 AND Any very high-power appliances

(b)

	TIER 0	TIER 1	TIER 2	TIER 3	TIER 4	TIER 5
Annual consumption levels, in kWhs		≥4.5	≥73	≥365	≥1,250	≥3,000
Daily consumption levels, in Whs		≥12	≥200	≥1,000	≥3,425	≥8,219

(c)

Figure 2-3: SEforAll multi-tier matrices for (a) household electricity supply, (b) household electricity services, and (c) household electricity consumption. The rating for a given household is equal to the lowest tier-rating across all attributes. Figure source: (Bhatia and Angelou, 2015)

2.5.1 Demand Density and Supply Costs

Because extending the grid in densely populated areas near to power infrastructure entails few new physical components, the associated marginal costs of extension to new consumers are relatively small. It wouldn't make sense to invest in largely redundant systems with high fixed costs covering the same areas. As such, the provision of electric power in dense demand areas is commonly considered a natural monopoly.

These marginal costs increase when the density of demand decreases. New lines of longer lengths become necessary, and relatively fewer new customers can benefit from such extensions. Compounding these difficulties, rural customers generally have lower demand for electricity than urban ones. This results in lower marginal revenues relative to costs.

In remote regions with lower population densities, off-grid technologies are generally more attractive choices for electrification relative to grid extensions. Mini-grids can lower the marginal cost of connection while still providing considerable supply in remote regions with an adequate density of buildings. Isolated units, such as solar home systems, may be most cost-effective in remote regions with lower building densities. These options exemplify sequentially lower levels of connectedness. Though their costs are higher with regards to generation, they benefit from larger savings on the capital expenses associated with interconnection.

2.5.2 Consumer Demand and Diminishing Marginal Utility

Consumer demand is highly important when considering the economics of electricity access. Understanding what people are willing to pay for electricity at a given level of service is a driving factor in understanding the quantities of electricity consumed and ultimately the revenues that planners can expect to recover.

A demand curve for a characteristic consumer is shown in Fig. 2-4, representing diminishing marginal utility with specific attributes. Customers are willing to pay high prices for the first units of electric energy serviced. These first units correspond to the highest utility uses of electricity, perhaps related to electric lighting, mobile device charging, or fan-based cooling. Their willingness to pay for additional consumption decreases precipitously, however, as subsequent uses are less essential and may more aptly correspond to luxuries than necessities.

It is important to note that individual consumers likely have varying preferences; Fig. 2-4 reflects the general nature of the demand curve one would expect consumers to have. This curve is also a function of the quality or reliability of the electricity service itself. In general, reliable service is more valuable than unreliable service as it can provide utility at times of highest need. Lastly, it is important to realize that these curves may evolve over time for different consumers, corresponding to changing affordabilities, technologies, the presence of potential substitutes, behaviors, economic conditions, and other factors (Pérez-Arriaga, 2017).

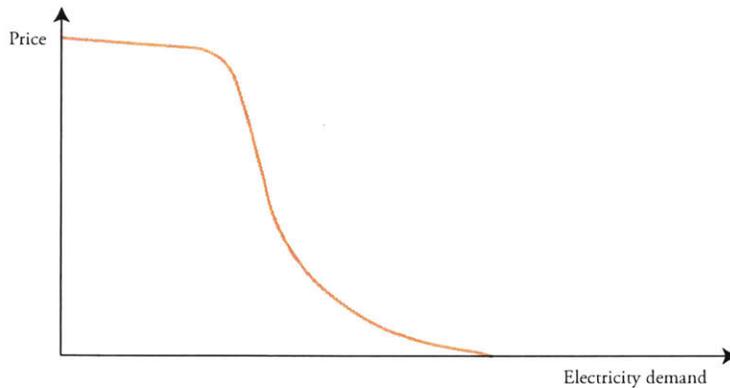


Figure 2-4: A characteristic demand curve for electricity. Figure source: (Pérez-Arriaga, 2017)

2.5.3 Subsidization and the Viability Gap

A central principle emerges from the economics of electricity access: rural electrification needs subsidies. It is most economical to provide electricity service of a given quality to customers in highly populated regions and most expensive to provide service to those in rural areas. For this reason and because distribution companies tend to apply uniform tariffs, subsidies are given to supply rural consumers with power in virtually every country in the world.

Rural electrification is enabled by cross subsidization through tariffs or with explicit public funding. Cross subsidization reflects the act of charging urban consumers tariffs that are higher than their supply costs, and passing profits on to rural consumers who are themselves serviced at a loss. Public funding, generated through taxes and other vehicles, is also used to facilitate rural electrification. Though subsidization is endemic to developed and developing countries alike, developing countries with low electrification rates face particular challenges due to the relative magnitudes of the subsidies required. The costs required to build new infrastructure and service large numbers of non-electrified consumers are generally far too high for cross subsidization and public budgets to cover, and as a result, large populations without access to electricity persist.

A useful concept when considering how to improve the economics of electricity access is the ‘viability gap.’ The viability gap is the difference between the costs of providing some quality of service and consumers’ willingness to pay for it (Pérez-Arriaga, 2017). In theory, this gap is present for nearly all non-electrified consumers. If it were not, entrepreneurs with access to financing would be able to make a profit off of providing electricity services to these populations. Because this is generally not the case (e.g., in equilibrium, ignoring distorting effects concerning crime and the destruction of property, etc.), one can think of endeavors towards universal electricity access as finding ways to close this gap. This can be done through the development and use of lower-cost technologies (e.g., solar panels, batteries, etc.), the sourcing of donor funding, the development of more suitable business models for a particular

context, and the design of plans making improved use of subsidies. It can also be done by promoting productive uses of energy and fostering economic growth from the bottom-up.

One avenue that is commonly pursued for profitably providing energy access is the exploitation of the high shoulder that is characteristic of consumers' demand curves for electric power, as shown in Fig. 2-4. Because consumers have their highest willingness to pay at low levels of demand, system designs that provide only very basic service are oftentimes viable. Basic mini-grids and stand-alone power systems are often able to serve this market; however, they are generally more limiting in that they have low ceilings to the levels of access they can provide. Pursuing this strategy oftentimes limits the types of productive services enabled.

2.6 Political Factors and Consumer Preferences

Above and beyond the economics of electrification, political factors and consumer preferences play major roles with regards to electrification. Regulators, multiple central and state agencies, utilities, developers, NGOs, development organizations, investors, and end consumers constitute a group of stakeholders that comprise multiple interrelated layers of decision-making. Agents often do only what is in their self-interests, and in aggregate decisions may be made that run counter to the interests of the end beneficiaries. Lindblom referred to such processes of how policy and institutional solutions oftentimes come to be as, "muddling through." He describes these processes as being highly incremental and composed of numerous individual elements (Lindblom, 2018). In many countries and in many cases, endeavors towards electrification can be characterized as muddling processes.

When choosing between supply technology options, one particularly salient consideration concerns the evaluation of consumer preferences. Regardless of whether the economics of a given system make sense, having sufficient customer buy-in will yield improved system maintenance, use, and ultimately cost-recovery and long term viability. Without such support, even the best projects may fail to become viable. For instance, if customers perceive off-grid systems to be intrinsically inferior to grid connection, they may consume less electricity and servicing an area may soon become inviable. Consumer participation in granular electrification decisions, education, and training may help in these cases.

2.6.1 Politicization

The politicization of electricity access has been shown to shape electrification endeavors. In his 2015 book, "Power and the Vote: Elections and Electricity in the Developing World," Brian Min provides a detailed account for how political manipulation shapes the distribution of electricity in India and other parts of the world. Min uses analyses based on nighttime lights satellite imagery to discern how democratically elected politicians have targeted electricity services. Among other findings, Min asserts that the competitive elections that characterize democratic governments

yield increased provisions of public services such as electricity than nondemocratic ones; democratic governments also do a better job at targeting electricity supplies to the poor. Additionally, Min presents studies for the Indian state of Uttar Pradesh that show how villages represented by winning political parties were supplied more power after elections than villages represented by losing political parties. Min’s hypothesis is that political actors provided and took away power to and from regions for strategic reasons (Min, 2015). This direction of inquiry demonstrates how political considerations can easily take precedence over solely efficiency-based ones in this field.

In some cases, politicization of electricity access has been perceived to distinctly harm the power sector. In some developing countries, heavily subsidizing tariffs is used to rally political support for decision-makers. One major byproduct, however, is that doing so exacerbates net operating losses associated with electricity service provision. Taken at scale, these dynamics can work to bankrupt distribution companies and provide disincentives for further rural electrification and reliable service provision.

2.6.2 Corruption

Corruption is a significant problem that affects infrastructure and economic development generally. While treating the intricacies of corruption is outside of the scope of this thesis, it should be noted that corruption can be a major impediment to rural electrification. It can stymie flows of donor funding, financing, and subsidy disbursement. In a recent study on infrastructure and urban development, the World Economic Forum reported that corruption on publicly funded construction projects can increase financial costs from 10-30% (Forum, 2017). In highly resource constrained settings, costs of these magnitudes are significant. Efforts to root out corruption have the potential to significantly benefit the provision of public services like electric power. Andrews et al. make the observation that too often, individual agents are blamed for instances of corruption observable by the public. Instead, the authors espouse a systemic perspective on the origins of corruption and propose organizational interventions including civil service, judicial, and public finance reform (Andrews et al., 2013).

2.7 Regulation and Business Models

Sound regulation and the selection of suitable business models are critical aspects to expanding electricity access. Electrification requires significant investment into physical infrastructure. Matching financing, capital and operating costs, and consumer affordability considerations are essential to viable electricity service provision.

The status quo business model is the “utility model,” or equivalently the “traditional distribution model” or a “regulated monopoly with a territorial franchise.” This would typically look like a publicly or privately owned utility subject to cost-of-service regulation; however, numerous variants exist. In-depth descriptions of the utility model are provided in (Pérez-Arriaga, 2014, 2017). In many contexts, the traditional model has proven inadequate for providing electricity access rapidly enough

to meet societal goals. Challenges include tendencies towards over-subsidization, utilities being in poor financial health, and corruption as discussed in Sections 2.6.1 and 2.6.2.

Alternative regulatory frameworks and business models are also being tested, which may be more suitable for off-grid consumers in the short-term. These include (1) unregulated entrepreneurs, (2) regulated entrepreneurs, and (3) licensed non-incumbent utility. Furthermore, new frameworks are being proposed, like Pérez-Arriaga et al.'s "Electricity Company of the Future." Those interested can refer to (Pérez-Arriaga et al., 2018) for more information.

2.8 Conclusions

Universal electricity access has received increasing levels of attention owing to its status as a key enabler of development. Economic and ethical reasons abound for the promotion of electricity access as a social goal. Nevertheless, considerations for and the implications of electrification are highly complex: political forces, evolving supply technologies, regulatory frameworks and business models, financing, consumer preferences, climate considerations, and others interact yielding a dynamic objective that is difficult to capture in its totality. Indeed, it is challenging to even define what electricity access means. As will be discussed in Chapter 4, there are concrete technical components in the electrification planning problem for which techno-economic models can optimize for; however, this chapter discusses numerous considerations for which, as of now, human judgment is required to capture. By providing a brief overview of important components to electricity access planning, this chapter aims to set the stage for more in-depth discussions about how planners, decision-makers, and citizens can more productively work towards universal electricity access goals going forward.

Chapter 3

Building Footprint Extraction and Localization

Remote sensing refers to the procurement of information about objects or phenomena without physical contact. It commonly refers to the use of aerial and satellite sensors to identify objects on Earth and determine their properties. In many cases, remote sensing allows for faster and more economical mapping efforts relative to on-site observation. The remote sensing literature describes two approaches to extracting information from image data: photointerpretation and quantitative analysis. Photointerpretation refers to the extraction of information through visual inspection of imagery by a human analyst. Quantitative analysis, on the other hand, involves a computer making geospatial judgments. The process of assigning labels to image pixels is a form of classification (Richards and Jia, 2006).

While human-based methods for image classification can confer high accuracies, automated computer-based methods have improved significantly with the development of modern machine learning techniques. Computer-based methods are highly advantageous for classification tasks at scales that would be prohibitively expensive if done by hand.

Because mapping geographic features and existing infrastructure is an integral part of infrastructure planning, remote sensing can be very valuable to planners, policymakers, and other stakeholders involved in these efforts. For the purpose of electrification planning in developing countries, mapping building locations is of particular importance. Planners need to know where buildings are located in order to plan on how to provide them with electricity. They also need to make determinations on whether these buildings are already serviced with power and what their demand characteristics will be over time.

The largest and most notable manual building labeling endeavor is the OpenStreetMap (OSM) project. OSM provides free and open detailed building and street annotations using a crowdsourcing-based approach: millions of participants conduct ground-based surveys and perform manual labeling on top of aerial imagery (OpenStreetMap, 2017). While OSM's data is impressive and rivals proprietary sources in terms of size and granularity, the quality of its data is inconsistent (Yuan, 2016). The availability of OSM's data is limited in developing countries, and this is espe-

cially true in rural areas. To procure complete building data sets in these regions for large areas without performing resource intensive surveys or manual labeling, automatic methods are required. As such, this chapter is devoted to a discussion of modern quantitative analysis techniques for classifying buildings in satellite imagery. Treatments of electrification status and demand growth will be presented in subsequent chapters. Specific tasks considered here include building footprint extraction and building localization. Building footprint extraction represents the endeavor of classifying pixels in image data that correspond to buildings. While most methods for this building extraction task are able to classify building areas, they are not able to discern individual buildings from other adjacent or nearby buildings. The task of building localization involves the next step of identifying individual building locations within clusters of adjacent pixels classified as belonging to buildings.

We will first provide background on building extraction methods based on traditional machine learning approaches and those based on deep learning approaches. We then describe training data procurement experiments and a building extraction system based on a popular convolutional neural network architecture for semantic segmentation. We present image quality and generalization studies using this model. Finally, we introduce building localization approaches we've developed and show how these techniques can be applied to identify buildings in large regions of India and Uganda.

3.1 Background

Machine learning approaches applied to the task of building extraction from satellite imagery can be partitioned into two groups: those using traditional computer vision and machine learning methods and those using deep learning-based methods. In this section, we provide brief overviews of neural networks, convolutional neural networks, and machine learning methods used for building extraction. The goal of this section is to provide context for sections on deep learning-based building footprint extraction and building localization to follow. For a comprehensive treatment of deep learning, (Goodfellow et al., 2016) is an excellent reference. Likewise, (Mnih, 2013) provides fairly comprehensive discussion on the topic of machine learning on satellite imagery.

3.1.1 Feedforward Neural Networks

Feedforward neural networks, also referred to as deep feedforward networks and multilayer perceptrons, have enabled the design of powerful models for classification. Their main advantage over traditional methods is that they are able to learn specific features from input data, rather than rely on hand-designed features. Feedforward neural networks are composed of layers of units, each carrying out simple computations. Fig. 3-1 shows a simple neural network architecture with an input layer, a hidden layer, and an output layer. Following the notation used in (Bishop, 2006), input variables are written as x_1, \dots, x_D ; they comprise the input layer and store the components of an input vector. Linear combinations of the input variables define ac-

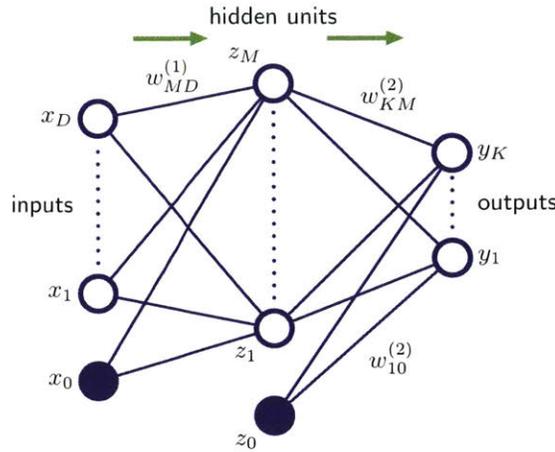


Figure 3-1: Simple feedforward neural network. Figure source: (Bishop, 2006).

tivations, a_j , where $j = 1, \dots, M$ as defined in 3.1. Parameters $w_{ji}^{(1)}$ represent weights and $w_{j0}^{(1)}$ represent to biases.

$$a_j = \sum_{i=1}^D w_{ji}^{(1)} x_i + w_{j0}^{(1)} \quad (3.1)$$

Each activation is transformed using a differentiable or subdifferentiable nonlinear activation function, $h(\cdot)$, to confer outputs of hidden units, z_j , within the hidden layer as shown in 3.2.

$$z_j = h(a_j) \quad (3.2)$$

The activation function is commonly defined to be the logit sigmoid 4.10, the hyperbolic tangent function, or the rectifier function, $\max(0, z)$.

$$\text{sigmoid}_{\text{logit}}(a) = \frac{1}{1 + \exp(-a)} \quad (3.3)$$

Linear combinations of outputs z_j from the hidden layer are subsequently used to define K output unit activations a_k where $k = 1, \dots, K$ as illustrated in 3.4.

$$a_k = \sum_{j=1}^M w_{kj}^{(2)} z_j + w_{k0}^{(2)} \quad (3.4)$$

Finally, output unit activation functions are applied to output unit activations to describe classification probabilities. The logit sigmoid function and the softmax activation function, defined by 3.5, are frequently used and for binary and multiclass classification problems, respectively.

$$\text{softmax}(a_k) = \frac{\exp(a_k)}{\sum_j \exp(a_j)} \quad (3.5)$$

The simple network described and presented in Fig. 3-1 can be generalized by adding additional hidden layers, changing activation functions, and modifying overall architectures. One generalization we use involves skip layer connections, which effectively define unit connections that skip a layer or layers in the neural network, as shown in Fig. 3-2. This example also demonstrates model sparsity: consecutive neighbors need not be fully connected.

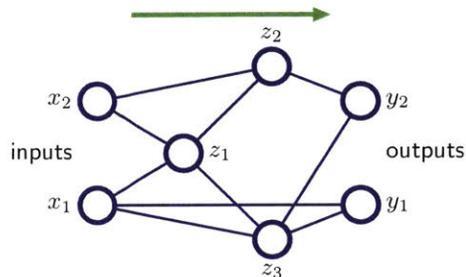


Figure 3-2: Sparse feedforward neural network with skip layer connections. Figure source: (Bishop, 2006).

At a high level, model parameters (weights and biases) for neural networks may be trained by defining an error function $E(\mathbf{w})$, randomly initializing parameter values, using a method called error backpropagation to evaluate the gradient of the error function, and employing iterative optimization methods like stochastic gradient descent to converge on a solution. This describes a procedure where information is sent forwards and backwards through the network and tune parameters in order to minimize the error function on training data. A general description of the error backpropagation algorithm is given in Algorithm 1.

Algorithm 1 Error backpropagation algorithm

- 1: Input: set input variables as training vector \mathbf{x}_n
 - 2: Feedforward: use $a_j = \sum_i w_{ji}z_i$ and $z_j = h(a_j)$ to forward propagate through the network
 - 3: Output error: evaluate δ_k for all output units using $\delta_k = y_k - t_k$
 - 4: Backpropagate error: backpropagate δ values for each hidden unit in the network using $\delta_j = h'(a_j) \sum_k w_{kj}\delta_k$
 - 5: Output: Compute the required derivatives using $\frac{\delta E_n}{\delta w_{ji}} = \delta_j z_i$
-

3.1.2 Convolutional Neural Networks

In the previous section, we discuss feedforward neural networks and present model architectures with fully connected layers. Generalizations are also mentioned that involve skip architectures and model sparsity. One such generalization, convolutional neural networks (ConvNet), represent a class of feedforward neural networks that have commonly been applied to image data; they generally employ local receptive fields,

weight sharing, and sub-sampling to effectively deal with image data. As shown in Fig. 3-3, the input layer is structured as a multidimensional array. For common imagery, the array is typically of size $h \times w \times d$, representing an image's height, width, and channel dimensions. Convolutional layers have units that are organized into planes called feature maps. Units within a given feature map take in inputs from small subregions or patches of the input image. In addition, these units are constrained to share weight values with all of the other units within the feature map. The use of local receptive fields allows for the exploitation of spatial structure within an image and renders the evaluation of activations within hidden units equivalent to performing convolution with a kernel comprising weight values. The use of weight sharing results in significant computational savings for deep networks by decreasing the number of parameters. It also confers approximate invariance to input image translations and distortions as units within a feature map effectively detect the same patterns in different parts of the input image.

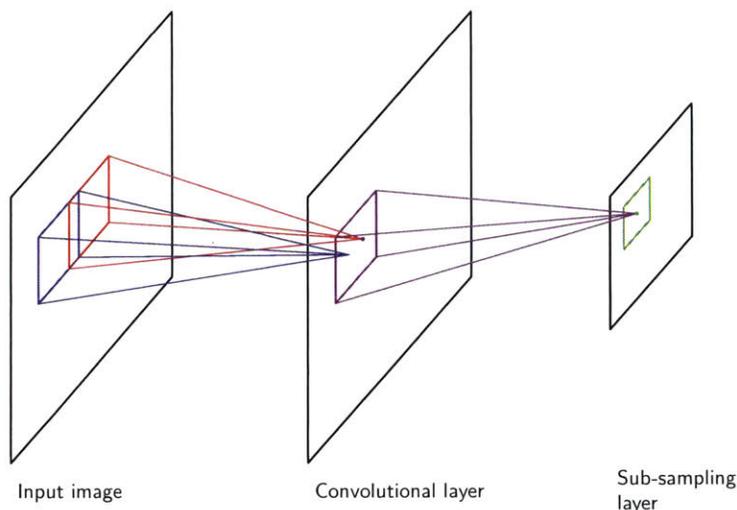


Figure 3-3: Layers in convolutional neural networks. Figure source: (Bishop, 2006).

Feature maps within a convolutional layer are typically paired with planes of units within a subsequent sub-sampling layer, also shown in Fig. 3-3. Units within the sub-sampling layer take in activations of convolutional layer units within small corresponding receptive fields. They then perform sub-sampling using functions like max pooling, which output patch-wise maximums. These values may then be scaled by adaptive weight and bias parameters and transformed using an activation function. Sub-sampling effectively reduces the spatial dimensions of the layer and the corresponding complexity of the features being captured.

Most contemporary ConvNets for image classification employ multiple feature maps within a given convolutional layer to detect different features at a given scale, and include several convolutional and sub-sampling layers. Feature maps within layers that are closer to the input image tend to learn lower-level features like edge detectors that describe local information. Deeper layers within a ConvNet aggregate lower-

level features and describe more semantically rich global information. For models that perform whole-image classification, the last layer or layers of a network are typically fully connected and like the feedforward networks described in Section 3.1.1, they employ logistic sigmoid or softmax activation functions for binary or multiclass classification.

3.1.3 Traditional Methods for Automatic Satellite and Aerial Image Labeling

A variety of traditional computer vision and machine learning methods have been applied to aerial and satellite image labeling, dating back to the late 1960s. The earliest systems used for automated labeling focused on the classification of broad classes of terrain using low resolution imagery, and were predominately knowledge-based approaches (Idelsohn, 1970; Bajcsy and Tavakoli, 1976). Machine learning-based approaches followed. Some notable examples employed Bayes' classifiers and neural networks. Neural networks were seen as advantageous over Bayes' classifiers because they could learn non-linear decision boundaries, whereas Bayes' classifiers typically employed multivariate normals for their class-conditional distributions and as such were limited to linear or quadratic decision boundaries (Decatur, 1989; Paola and Schowengerdt, 1995; Mnih, 2013). Neural networks were also modified to learn local contextual features using small pixel windows and hand-designed features (Bischof et al., 1992; Boggess, 1993; Haralick, 1976).

With the increased availability of higher resolution satellite imagery in the early 2000s, smaller objects including buildings and roads could be discerned and classified. The Ikonos and Quickbird satellites, for example, started producing panchromatic images at ~ 1 meter resolutions in 1999 and 2001, respectively. As a result, different techniques were employed over time that could account for increasingly more complex features. They typically relied on filters from popular filter banks for computing input representations (Mnih, 2013). Support vector machines (SVMs) were applied to pixel classification tasks using low-level features, and were seen as advantageous over neural networks due to their convex loss function and resultant resistance to local optima (Huang et al., 2002; Song et al., 2005; Mnih, 2013). Ensemble methods were also commonly applied to high resolution imagery. Porway et al., Dollar et al., and Nguyen et al., developed notable models based on boosting to detect objects such as cars, roads, and buildings (Dollar et al., 2006; Nguyen et al., 2007; Porway et al., 2008). Kluckner et al. used random forests for the semantic classification of buildings, streets, green areas, trees, and waterbodies as well (Kluckner et al., 2009; Kluckner and Bischof, 2009). Because of their robustness to mislabeled data and their ease of being parallelized, random forests methods are thought to be more appropriate to aerial and satellite image classification than typical boosting methods (Mnih, 2013).

Segmentation and superpixel classification approaches were used as an alternative to the classification of individual pixels. In these approaches, images are first segmented into superpixels: regions of similar pixels in an image. Statistical approaches are then used to classify these superpixels. He et al. demonstrate this methodology

using conditional random fields on non-aerial imagery, Huang and Zhang use segmentation approaches and SVMs to extract road centrelines, and Varshney et al. use segmentation and random forests classifiers to classify building footprints (He et al., 2006; Huang and Zhang, 2009; Varshney et al., 2015). The purpose of using segmentation is to reduce the computational complexity of image labeling, as there are by definition fewer superpixels than pixels for a given image (Mnih, 2013).

3.1.4 Convolutional Neural Networks for Automatic Satellite and Aerial Image Labeling

Ever since the convolutional neural network AlexNet won the ImageNet competition in 2012 with a 10.8 percentage point margin, ConvNets have exploded in popularity for computer vision tasks. They have often proven more effective than other methods for computer vision problems (MIT Technology Review, 2014). Aerial and satellite image labeling is no exception. Mnih et al. apply ConvNet architectures in 2010 to perform post-classification on neural network outputs for road detection (Mnih and Hinton, 2010). They also use ConvNets with untied weights and robust loss functions for high performance extraction (Mnih and Hinton, 2012; Mnih, 2013). Among subsequent academic studies, Yuan’s 2016 work on using ConvNets for building extraction from aerial scenes stands out. Yuan uses a signed distance function from building boundaries in output representation for ConvNets, and shows how this can benefit classification performance and enable the representation of fine-grained labels for border boundaries (Yuan, 2016). Zhang et al. discuss Facebook’s efforts to generalize these methods to perform building detection on a global scale with 500TB of imagery (Zhang et al., 2017).

Facebook’s efforts are indicative of greater industry interest in object extraction from satellite imagery. Facebook is working on extracting building footprints from satellite imagery to help inform its Internet access efforts. In 2016, the company announced that it will release a 5 meter resolution data set for 20 countries around the world in partnership with the World Bank and Columbia University’s Center for International Earth Science Information Network using DigitalGlobe’s high-resolution satellite imagery (Gros and Tiecke, 2016; Tiecke, 2016). At the time of writing, population estimates for 24 countries have been released on the Center for International Earth Science Information Network’s High Resolution Settlement Layer website; however, they are only available at ~ 30 meter resolutions (Columbia University Center for International Earth Science Information Network, 2018). In addition, though Facebook’s access to high resolution DigitalGlobe imagery affords the possibility to release building-level GIS data, the population estimates publicly released in the High Resolution Settlement Layer data sets are derived from aggregate census metrics and may lack levels of precision desired for many infrastructure planning endeavors. In essence, the 30 meter resolution pixels provided can be interpreted as binary classifications for whether buildings exist in the corresponding areas. Population estimates for these pixels are derived from region-level aggregate data, not corresponding satellite imagery.

A number of other companies are doing feature extraction from high resolution satellite imagery as well using ConvNets. Toronto-based Ecopia Tech Corporation and Mountain View-based Orbital Insight are extracting a variety of objects from high resolution satellite imagery. They are both working on top of DigitalGlobe GBDX, DigitalGlobe’s cloud-based platform for accessing their imagery (Ecopia, 2017; DigitalGlobe, 2017; Orbital Insight, 2017).

3.1.5 Satellite imagery

The satellite image industry is changing in significant ways, and is expected enable cheaper and more accurate remote sensing going forward. Earth imagery is becoming more plentiful, with greater spatial resolution, temporal resolution, and coverage year after year. DigitalGlobe announced the availability of very high resolution 30 cm satellite imagery products with its WorldView-3 and WorldView-4 satellites in 2015 and 2017 (Van Uum, 2015; Ray, 2017). In 2017, DigitalGlobe competitor Planet, Inc. acquired Terra Bella and its high resolution satellites from Google. Planet also launched 88 miniaturized satellites on February 14, 2017; as of that date, the company operated 149 satellites and achieved the scale to image all of the Earth’s landmass at medium-resolution every day (Marshall, 2017; Robbie, 2017).

The utility of these satellite products for infrastructure planning and electrification planning in particular are a function of several variables including spatial resolution, temporal resolution, and cloud coverage, spectral range, and off-nadir angle. Higher spatial resolution allows for better building and road extraction accuracies and higher temporal resolution allows for the attainment of more recent imagery and improved confidence in infrastructure representations for planning. Low cloud coverage enables improved representations as well, and reduces the need for acquiring multiple images over the same areas. Improved spectral range may help classification tasks such as the delineation of foliage from man-made infrastructure. Finally, having a higher off-nadir angle may help to improve visibility of the sides of buildings (Varshney et al., 2015).

Paid imagery and free imagery are commonly available for most places in the world. The Google Maps API allows for free downloading of visible-light images. Imagery can also be purchased from DigitalGlobe and its third party resellers. Both approaches were used for the studies presented in this thesis. WorldView-2 satellite imagery was purchased from Apollo Mapping, a distributor for DigitalGlobe, and corresponding imagery was downloaded using the Google Maps API. These data sets will be elaborated on Section 3.3.1.

3.1.6 Convolutional Neural Networks for Semantic Segmentation

In 2014, Long and Shelhamer et al. presented the novel approach of Fully Convolutional Networks (FCNs) for semantic segmentation. At the time, this approach represented the state-of-the-art in semantic segmentation and has since set a standard for continued improvement. FCNs for semantic segmentation are trained end-to-end,

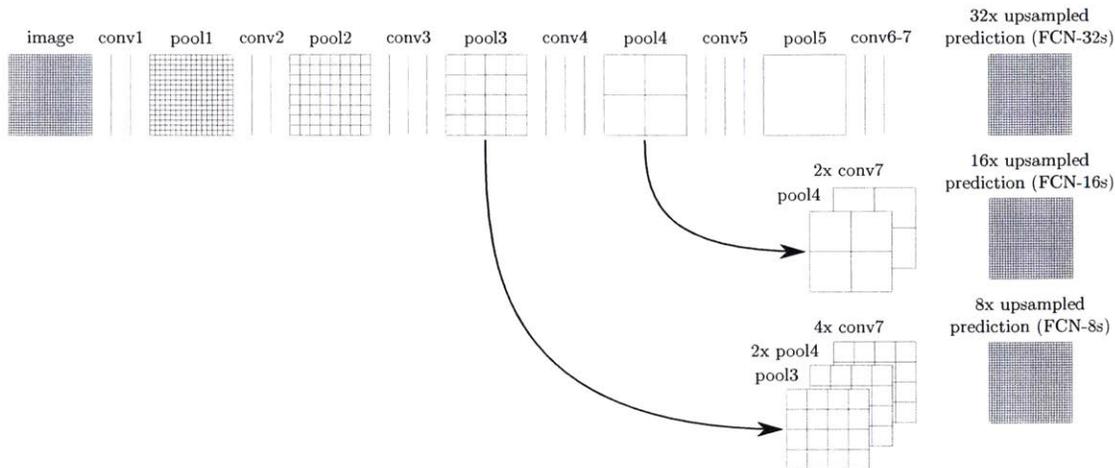


Figure 3-4: Model architectures for fully convolutional networks for semantic segmentation. Figure source: (Long et al., 2015).

provide pixel-to-pixel prediction, and can employ supervised pre-training. They use deconvolution layers for upsampling and take advantage of skip architectures to combine semantic and appearance information from multiple layers contained within the network. Fig. 3-4 shows model architectures for $32\times$, $16\times$, and $8\times$ upsampling based on combining predictions from intermediate layers within the network. The authors report that the $8\times$ upsampled model produce the best results, and demonstrate 62.2% mean pixel intersection over union (IU) scores on the PASCAL VOC 2011 dataset, which has 21 classes including the background. The work by Long and Shelhamer et al. builds off of the concept of FCNs pioneered by Matan et al. and the concept of jets pioneered by Koenderink and Van Doorn (Long et al., 2015). In 1991, Matan et al. used FCNs to recognize whole series of digits as opposed to digits individually (Matan et al., 1992). In 1987, Koenderink and Van Doorn used local jets to give rich representations of geometry and semantics with filters on multiple scales (Koenderink and van Doorn, 1987).

Since the work of Long and Shelhamer et al., several other methods have been explored to improve the performance of convolutional neural networks for semantic segmentation. In 2015, Chen and Papandreou incorporated probabilistic graphical models in the form of fully conditional random fields (CRF) to overcome poor localization (Liang-Chieh et al., 2015). Later in 2015, Zheng and Jayasumana showed that unpacking dense CRFs into individual computations and joining them to the network yields further improvement (Zheng et al., 2015). In 2015, Noh et al. demonstrated a novel semantic segmentation algorithm that incorporates a learned deconvolution network for even better performance (Noh et al., 2015). Finally, in 2016, Kendall et al. introduce Bayesian SegNet for semantic segmentation. They predict pixel-wise class labels with a measure of model uncertainty from combining Monte Carlo sampling with dropout at test time. Such representations of model uncertainty have been shown to improve classification performance, especially for smaller datasets (Kendall et al., 2016).

3.2 Procuring Training Data

The procurement of training data for supervised machine learning is often a time and resource-intensive endeavor. We find the compilation of training data for building extraction to be no different. The models we use for the semantic segmentation require segmentation training data where annotations are captured at the pixel level. While the OSM data set provides manual labels that may seem promising for model training, its sparsity, inconsistency, and the effort of aligning it with specific satellite images at the pixel-level render it less useful for our application in developing countries. Due to these considerations and the fact that our building extraction models may not be highly generalizable across regions, we found it beneficial to procure our own training data for regions of interest.

Even when controlling many parts of the annotation process, we generally found that it was difficult to produce consistent training sets across multiple annotators. Tracing building outlines in satellite imagery is more subjective than one may think. It is also tedious work and requires a level of patience that is not immediately rewarded. We generally found that annotators with less stake in the results of the automatic building extraction task were less likely to produce high quality annotations. This finding is validated in the literature. Zhang et al. report that their labelers were only 85% accurate at labeling buildings from satellite images (Zhang et al., 2017).

3.2.1 Using Microsoft Paint and Google Earth

We initially experimented with Microsoft Paint and Google Earth for labeling buildings on local computers. In Microsoft Paint, annotators use the polygon tool to color over buildings in satellite imagery. We then use scripts to create appropriately coded annotation masks by querying polygon colors. In Google Earth, annotators add polygons to create KMZ files corresponding to building rooftops. Microsoft Paint is convenient because curators have full control over input imagery and can ensure that annotators look at image tiles one-by-one. Conversely, annotators using Google Earth need to pan around and search for buildings themselves; this approach is more prone to annotators missing buildings in regions of interest. Nevertheless, using Google Earth is beneficial because the KMZ files generated are appropriately georeferenced and can be used to make training data for corresponding satellite images with different resolutions in a straightforward manner.

Despite the fact that Microsoft Paint and Google Earth provide familiar and easily accessible platforms for image annotation, we find that their biggest drawback is their confinement to local computers. Coordinating tasks among several annotators is difficult to do, as files must be distributed and stored locally. In addition, using these platforms poses barriers to crowdsourcing and contracting out annotation tasks. Sending out tasks and consuming labels requires file transfer between unfamiliar parties. This is a cumbersome process and can be perceived as a security risk from both sides. To overcome these issues, we modified and built web-based tools.

3.2.2 Amazon Mechanical Turk

Following in the footsteps of other training set compilation efforts for supervised machine learning and computer vision, we tried using web tools that are compatible with Amazon Mechanical Turk (MTurk) for distributing and managing annotation tasks (Su et al., 2012; Zhou et al., 2014). Amazon Mechanical Turk creates a market where developers or ‘requesters’ can match with workers and distribute tasks that require human intelligence using an API. The MTurk workforce is globally distributed and is large enough to constitute a 24×7 service which uses a reputation system to allow degrees of targeting. Requesters can design an MTurk compatible web template using HTML, CSS, and JavaScript and program it with appropriate quality tests; requesters then provide a description, specify assignment durations and rewards, set required worker approval rates, and finally stipulate whether or not workers need to have achieved "Master" status. A given task is called a Human Intelligence Task (HIT); the same HIT can be distributed to multiple workers as different assignments. Assignment attributes and HITs are tracked by the MTurk system, and requesters can choose whether they deem assignments as satisfactory before paying workers (Amazon Web Services, 2017).

TurkCleaner Tests on Amazon Mechanical Turk

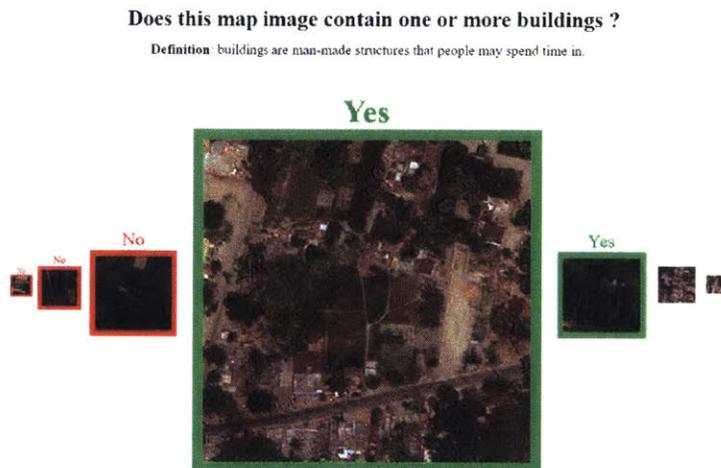


Figure 3-5: The TurkCleaner interface for the binary classification of whether or not buildings are shown in satellite image tiles.

The building extraction methodologies we will describe require satellite imagery to be divided into smaller image tiles in order to avoid memory issues during learning and inference. While image data obtained from the Google Maps API already comes in tiled form, purchased or otherwise procured satellite image strips need to be tiled while preserving information about their geographic extents. This tiling requirement also benefits the training data procurement pipeline. Annotators using web-based

tools only need to download the parts of satellite images that concern their annotation task.

Because the vast majority of the image tiles in our rural areas of interest do not contain buildings, we first use the TurkCleaner interface to do binary classifications of whether or not buildings are present in an image tile. The TurkCleaner interface is part of Dr. Jianxiong Xiao's open-source software library called the "Professor X Toolkit." It was formerly known as the Princeton Vision and Robotics Toolkit (Xiao, 2016). An image of TurkCleaner applied to our image tile classification task is shown in Fig. 3-5.

TurkCleaner allows users to quickly move through a set of images and label them with a "Yes" or "No" for a given query using their computer keyboards. The design shows workers previously annotated and upcoming images; this helps users to make sure they do not skip any images, assists them in anticipating upcoming tasks, and allows them to backtrack if necessary. The interface also enables requesters to incorporate built-in tests. The requester can manually label a few images and include them in their HITs. Workers will then label these built-in tests with the rest of their set without knowing which images constitute tests and which images correspond to their labeling contributions. The interface will reject any assignment submissions that do not meet requester-defined thresholds on the tests. By setting the threshold to a high-enough level, requesters can improve the likelihood that accepted work meets their quality requirements.

In one of my TurkCleaner experiments, I required users to classify whether or not image tiles had buildings in them for 5,280 image tiles corresponding to a 50cm resolution WorldView-2 image strip taken of a region in Bihar, India. I divided the task into 53 assignments, each with ~ 100 image tiles, and stipulated that while workers did not need to be MTurk "Masters," they must have a HIT approval rate of at least 85%. I also included 5 built-in tests for each assignment and required that at least 75% of these tests needed to be correct before workers could submit their jobs. Finally, I posted each assignment with a \$0.30 reward. Workers completed all 53 assignments in less than 3 hours with high quality classifications.

DrawMe Tests on Amazon Mechanical Turk

Once images are identified that have buildings in them, the next step is semantic training data labeling: drawing polygons around the buildings corresponding to their footprints. For this task, we modify the DrawMe project for semantic image labeling from the Professor X Toolkit (Xiao, 2016, 2017; Xiao et al., 2010). DrawMe allows for requesters to designate an image used for a given HIT. Unlike the TurkCleaner tool, however, DrawMe provides no clear way for integrating built-in tests. Workers draw polygons using their mouse and click to designate polygon vertices. Once at least one polygon is drawn, the user has the ability to submit their assignment.

Initial building labeling experiments on MTurk were done using a slightly modified version of DrawMe 1.0 with directions instructing workers to label all buildings in the images presented. In one of our experiments, we gave workers 2 minutes to trace building outlines, and awarded them workers \$0.50 if we accepted their assignment.

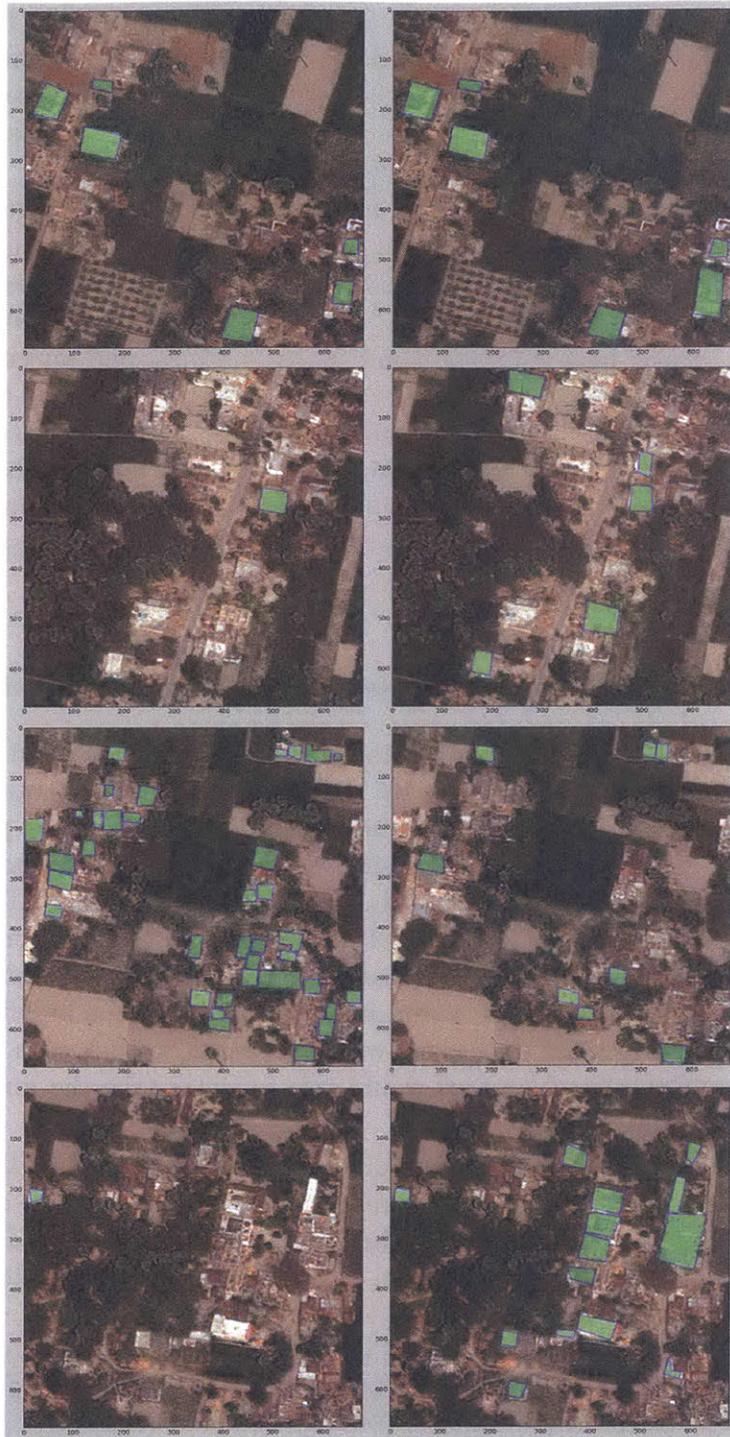


Figure 3-6: (left column) Sample annotations from workers who did not necessarily have to be "masters." (right column) Sample annotations from "master" workers. Neither populations provide high quality annotations using the DrawMe interface.

Fig. 3-6 shows a sample of our results. The images show building labels overlaid onto corresponding images using blue borders and green fill. The left column depicts annotations done by workers who did not necessarily have to be "masters." The right column depicts annotations for these same images but where workers needed to be "masters." Neither populations provided high quality annotations.

The problem with this interface for the task of comprehensive building footprint annotation is that workers can submit their assignment after just one building is annotated. Workers likely don't have an incentive to fully read or follow the instructions, and instead respond to the business logic portrayed by the user interface. Furthermore, different images have different numbers of buildings and varying annotation difficulties. Because there is no easy way to scale compensation with difficulty, even if workers fully annotated buildings according to the instructions, the high variance in assignment difficulties would incentivize completion of the easiest assignments and disincentivize completion of those that are more difficult. For these compounding reasons, it was clear that modification was necessary to source a higher quality training set.

DrawMe and Back End System Tests on Amazon Mechanical Turk

We modified the DrawMe tool and experimented with the approach of giving MTurk workers multiple image tiles per HIT. By storing intermediate annotation representations and asking workers to only annotate one additional building per image tile, we could ensure that they consistently did the minimum work necessary to improve our representations. We could also make the difficulty of various HITs much more consistent by specifying a fixed number of image tiles per HIT. If workers were presented with 20 image tiles, they would need to label 20 buildings in order to submit their job, and we would pay them commensurately. If all of the buildings in a given image tile were already labeled, workers were instructed to click on a toggle box denoting "Labeling Complete," and we would subsequently take the image tile out of circulation for annotations.

We built a back end system to coordinate the generation of HITs according to our latest representations, to collect completed HITs, and to update our annotation database. We used Python, NumPy, SQLAlchemy, SQLite, Boto, and aspects of the open-source simple-amt project (Johnson, 2016). A representation of our database schema is shown in Fig. 3-7.

While this approach seemed to be an improvement over DrawMe on MTurk alone, it still provided no way to ensure that building labels were done with high quality. Workers oftentimes did not cover building footprints tightly and also erroneously marked the "Labeling Complete" toggle box when buildings were still left unlabeled. We concluded that we may have been reaching the limits of the MTurk workforce's capabilities, and that the task of building labeling on satellite imagery containing multiple buildings may be too difficult without annotation training and effective two-way communication. Instead, we looked to freelancing platforms and personal relationships to provide the workforce necessary for completing our training sets.

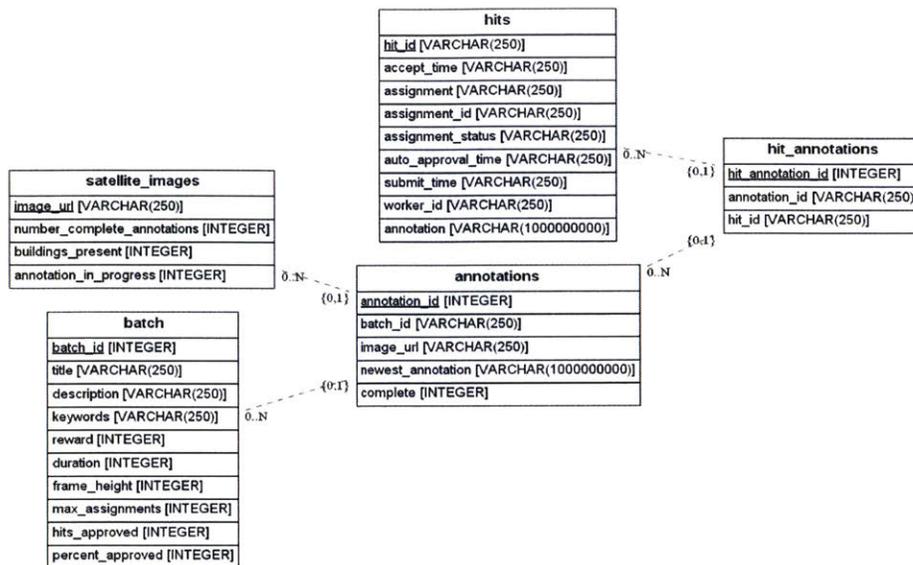


Figure 3-7: Database schema for MTurk back end tool. Though the tool improved annotation results, the annotations ultimately obtained were of lower quality than needed for our training sets.

3.2.3 Using Web Tools, Google’s Realtime API, freelancer.com, and Expert Annotators

Moving away from MTurk, we built interfaces to do manual image annotation employing workers with whom we could directly communicate and instruct. We replaced our back end system from the previous section with a system built on top of the Google Realtime API so that we could see workers’ progress in real time. The Realtime API makes it easy for workers to update and save data models while enabling requesters to view annotations as they are updated. It was also easy to integrate and did not require us to run a server to collect annotations.

A further modified version of the DrawMe interface is shown in Fig. 3-8. In this latest version, we used AngularJS to help to build a more intelligent and self-contained user interface. We allow workers to annotate images as before; however, they are also able to track their progress for a large set of images using a scroll down list. On the other end, requesters can monitor workers’ progress, sample the annotations, and help guide workers to fine-tune their annotations.

We tried working with freelancers on freelancer.com and personal contacts as ‘expert’ annotators. Workers from both populations provided satisfactory annotations; however, we ultimately found that it was easier to hire and work with local personnel. Freelancers oftentimes had slow Internet connections and experienced difficulty connecting to the Google Realtime API server. It was also difficult to coordinate with freelancers and enforce high quality annotations as we found that they oftentimes delegated their work to numerous employees working under them. Instead, working with local annotators generally led to better results with less effort on our end.

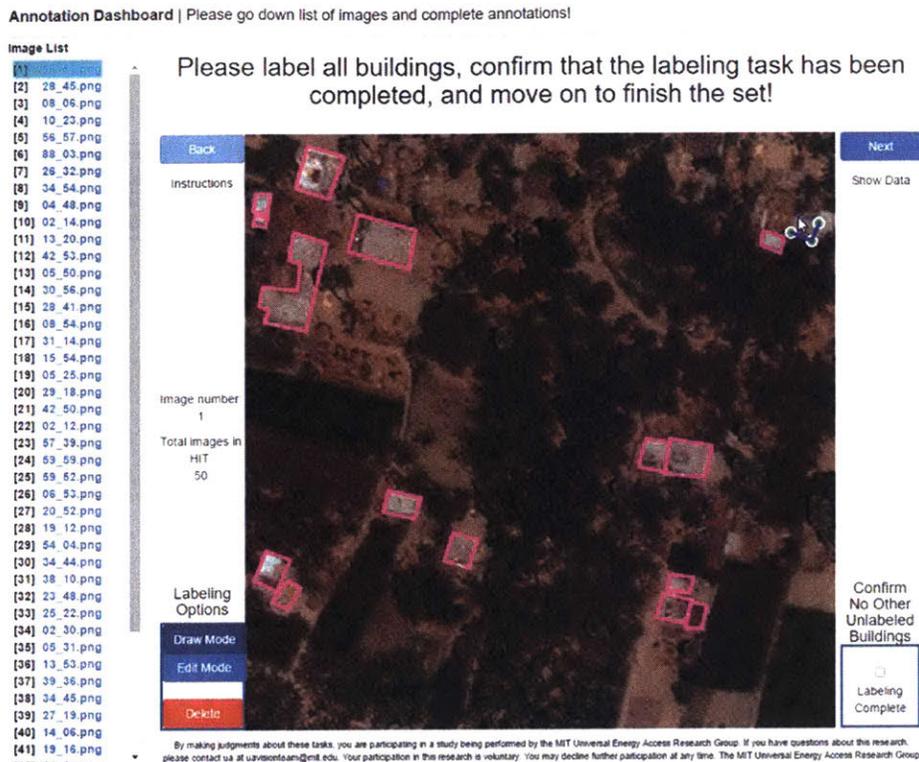


Figure 3-8: A modified DrawMe interface is shown. Our user interface presents annotators with a set of images and allows them to navigate the set and denote completion of annotations with ease.

3.3 Building Extraction and Image Quality Studies

In this section, we present an adaptation of Long and Shelhamer’s FCN models for semantic segmentation to the problem of building footprint extraction from satellite imagery (Long et al., 2015). While these models no longer represent the state-of-the-art methods for semantic segmentation in the computer vision community, we show that they are able to perform building rooftop extraction with satisfactory performance. Their relative simplicity also allows us to establish a basic foundation for continued research using convolutional neural networks for this application.

This section is divided into five parts. We first briefly discuss data procurement and formatting. We then review popular error metrics for semantic segmentation tasks and compare our FCN model results to previously reported methodologies. Finally, we provide an overview of image quality experiments and data augmentation experiments.

3.3.1 Data Procurement and Formatting

We obtained geospatial vector data in the ESRI Shapefile format from Varshney et al. corresponding to 596 buildings from 10 rural villages in Odisha, India. Varshney

et al. also shared a high-quality color balanced and orthorectified image of the area, which was purchased by MIT and originally taken by DigitalGlobe’s WorldView-2 Satellite at 50 cm resolution. The image has dimensions of 13,488 x 10,925 pixels, spanning just less than 37 km² (Varshney et al., 2015).

Because there were several villages and building structures in the image that did not have corresponding label information, we manually identified areas with missing labels to ignore in the training process. We then created raster masks that are pixel-mapped to the satellite image. A depiction of part of the mask which represents our ground truth is shown in Fig. 3-9. In the mask, black pixels denote background area, red pixels denote area covered by buildings, and tan pixels show areas which are either missing labels or correspond to ambiguous regions. The image regions corresponding to the tan areas are ignored during training. Large stand-alone examples of such regions correspond to our identification of areas that are missing labels while the borders of buildings are assessed to be ambiguous and rendered on the masks by dilating building shapes.



Figure 3-9: Raster masks as ground truth. Buildings are shown in red and regions that are ambiguous or missing labels are in tan.

The large image was broken up into south and north regions for training and testing, respectively. The training set comprises 55% of total land area, while the test set comprises 45%. The line of latitude for splitting the image was determined to avoid bisecting any of our 10 villages. As such, all villages are fully contained in an individual dataset. Finally, the images and ground truth masks were split into equally sized tiles by geographic coordinates and are approximately 500 x 500 pixels. Tiling the images helped to alleviate memory problems encountered when using the FCN models.

3.3.2 Error Metrics

The selection of appropriate error metrics was a topic of consideration due to the unique nature of our dataset and the application of rural electrification. Ignoring the tan unlabeled and ambiguous areas in the ground truth labels, 99.94% of the image corresponded to background area, while only 0.06% denoted building area. This makes it incredibly easy to achieve high metrics for guessing background pixels. Indeed, a rudimentary algorithm that labels all pixels as background would achieve

99.94% pixel accuracy. Because of this, other pixel-based metrics are important to consider in evaluating the success of models for our application. Common metrics for semantic segmentation include pixel accuracy, mean accuracy, IU, mean IU, and frequency weighted IU. Consistent with Long and Shelhamer, we define n_{ij} as the number of pixels of class i predicted to belong to class j , n_{cl} as the number of different classes, and $t_i = \sum_j n_{ij}$ as the total number of pixels of class i (Long et al., 2015). The metrics are given below:

$$\frac{\sum_i n_{ii}}{\sum_i t_i} \quad (3.6)$$

Pixel Accuracy

$$\frac{1}{n_{cl}} \frac{\sum_i n_{ii}}{t_i} \quad (3.7)$$

Mean Accuracy

$$\frac{\sum_i n_{ii}}{t_i + \sum_j n_{ji} - n_{ii}} = \frac{TP}{TP + FP + FN} \quad (3.8)$$

Intersection over Union (IU)

$$\frac{1}{n_{cl}} \frac{\sum_i n_{ii}}{t_i + \sum_j n_{ji} - n_{ii}} \quad (3.9)$$

Mean IU

$$\frac{1}{\sum_k t_k} \frac{\sum_i t_i n_{ii}}{t_i + \sum_j n_{ji} - n_{ii}} \quad (3.10)$$

Frequency Weighted IU

Because of the high disparity between the shares of background and building pixels, we note that the measures of pixel accuracy and frequency weighted IU are highly misleading. Both these measures become dominated by the assessment of background pixels and fail to assess the ability to accurately detect building pixels. The mean IU and the mean accuracy metrics are similar but differ in the denominator of the summation term. An interpretation of the difference is that the mean IU penalizes for false positives, while the mean accuracy does not. False positives are certainly a problem in our application: they yield situations where we plan for energy systems to incorporate buildings that do not actually exist. Because of this, we feel mean IU is a more appropriate metric for our application than mean accuracy. Finally, we compare mean IU to building IU and determine that, because of the expected high value of background IU and positive correlation between background IU and building IU, these measures are likely to both appropriately describe the relative merits of models but with different magnitudes. Since mean IU considers the success of background pixel classification in addition to building pixel classification, we feel it is the better choice for considering the quality of the overall model; however, building IU would have more merit if we were to expand our application to multiple

classes.

3.3.3 Assessment of FCN Semantic Segmentation Quality

We fine-tune FCN models for semantic segmentation based on a ImageNet VGG Very Deep 16 model incorporating Long and Shelhamer et al.’s skip architecture with lowest stride size of 8 pixels. We present results in Fig. 3-10 and Table 3.1. Fig. 3-10 shows sample building segmentations after fine-tuning and testing on WorldView-2 imagery; in addition, it shows segmentations after doing the same for Google Maps imagery. Table 3.1 shows pixel accuracy, mean accuracy, background IU, building IU, and mean IU metrics for image quality experiments (i.e., comparing WorldView-2 imagery to Google Maps), data augmentation experiments (i.e., comparing WorldView-2 imagery to an augmented WorldView-2 imagery dataset with horizontally flipped images), and model comparisons. Furthermore, FCN learning curves are depicted in Fig. 3-11. We discuss our image quality and data augmentation experiments in subsequent sections and present the model comparisons here.

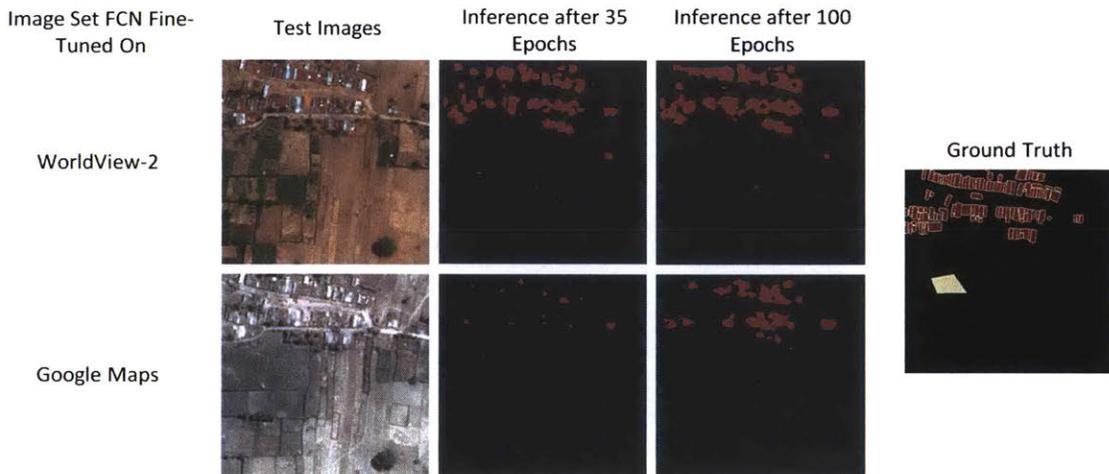


Figure 3-10: Inferences from ImageNet VGG Very Deep 16 models of the fcn8 type that were fine-tuned and tested on WorldView-2 images (top) and Google Maps (bottom) for 35 (second column) and 100 epochs (third column).

We quantitatively compare segmentation results from the FCN model to results from the Varshney et al. study’s seed-based region-growing image segmentation algorithm and random forest classifier (Varshney et al., 2015). We obtained these model results in collaboration with Prof. George Chen, a co-author of the Varshney et al. study. We show semantic segmentation results that outperform the Varshney et al. study’s segmentation and random forest approach. As shown in Table 3.1, after 100 epochs, the WorldView-2 fine-tuned models achieved 47.19% building IU and 73.58% mean IU while the model pipeline presented in Varshney et al. achieved 41.34% and 70.64%, respectively, for this same task. Part of the reason for this success likely results from the fact that FCN model learns richer features that more accurately rep-

resent the presence of buildings than the segmentation and random forest classifier approach. Varshney et al.’s classifier uses a feature vector composed of the region’s color distribution, area, diameter, ratio of perimeter to area, standard deviation of grayscale values, and a set of color ratios (Varshney et al., 2015). In contrast, the FCN model learns its own complex features at a more flexible range of scales.

The general strengths and weaknesses of the FCN model with regards to building extraction from satellite imagery can be assessed when qualitatively observing inference samples as shown in Fig. 3-10. While overall location and area information appears to be highly correlated with our ground truth images, it appears that the model has difficulty defining the exact borders of images. This can be due to the fact that our training data and ground truth labels are not precise about defining borders. The image data resolution makes exact borders somewhat ambiguous, resulting in the need for us to label these areas as regions to ignore in our ground truth. Furthermore, the model itself may have intrinsic difficulty detecting features for semantic segmentation at such small scales due to its lowest stride size of 8px.

Table 3.1: Semantic segmentation metrics for the FCN runs discussed and for the Varshney et al. comparison

Test Set	Epochs	Pixel Acc.	Mean Acc.	Background IU	Building IU	Mean IU
WorldView-2	35	99.95%	75.96%	99.95%	38.65%	69.30%
WorldView-2	100	99.96%	79.14%	99.96%	47.19%	73.58%
Google Maps	35	99.94%	52.01%	99.94%	3.98%	51.96%
Google Maps	100	99.95%	65.40%	99.95%	26.67%	63.31%
WorldView-2 w/ Flip	35	99.96%	78.74%	99.96%	45.29%	72.62%
WorldView-2 w/ Flip	100	99.96%	78.37%	99.96%	46.65%	73.30%
Varshney et al. on WorldView-2	N/A	not computed	not computed	99.95%	41.34%	70.64%

3.3.4 Image Quality

The comparison study between our FCN models fine-tuned and tested on WorldView-2 images and Google Maps images first required the procurement of Google Maps images. We used the Google Maps API to download image tiles from the region of interest and stitched them together into a seamless georeferenced image of 11,739 x 10,098 pixels. We processed the images and ground truth labels in the same way described previously. The new and standardized image tiles for the WorldView-2 images and Google Maps images were used to fine-tune FCN models for semantic segmentation based on a ImageNet VGG Very Deep 16 model incorporating Long and Shelhamer et al.’s skip architecture with lowest stride size of 8 pixels, as described previously.

Error metrics obtained from our comparison study are shown in Table 3.1, inference samples are shown in Fig. 3-10, and learning curves are depicted in Fig. 3-11. Our models learned features faster when fine-tuned with WorldView-2 data than with Google Maps images, as represented by steeper objective energy learning curves at each epoch. Our models fine-tuned with WorldView-2 data also showed more desirable

error metrics overall. After 100 cycles, the WorldView-2 fine-tuned models achieved 47.19% building IU and 73.58% mean IU, while the Google Maps fine-tuned models only achieved 26.67% and 63.31%, respectively. The differing qualities of detection can be seen when comparing inferences in the first and second rows of Fig. 3-10. The superior results obtained from fine-tuning with the WorldView-2 dataset owe credit to the higher resolution, improved color balance, and higher contrast of these images. The Google Maps images, in contrast, had faint “Google” watermarks, and minor and infrequent artifacts from degree-to-pixel rounding errors from the stitching process.

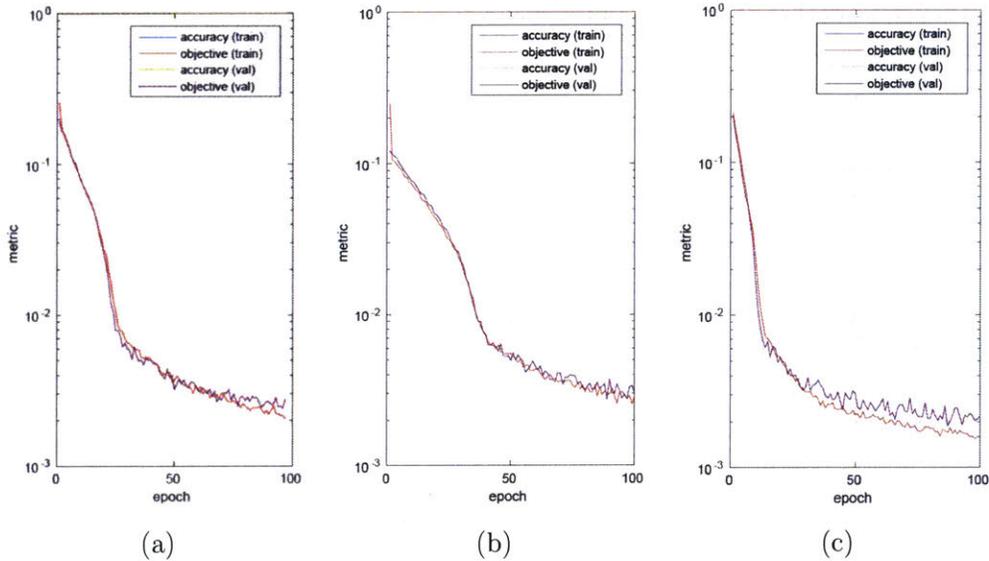


Figure 3-11: Learning curves shown for FCNs fine-tuned with (a) WorldView-2 images alone, (b) Google Maps images, (c) augmented WorldView-2 images.

3.3.5 Data Augmentation

In our data augmentation comparison study, we used the same FCNs for semantic segmentation, WorldView-2 images, and ground truth labels as described above; however, we effectively double the training set size by horizontally flipping our image tiles. With this change, we investigate the effects of data augmentation through flipping on performance.

We obtained interesting results from our comparison of networks fine-tuned with WorldView-2 images and networks fine-tuned with this same dataset but augmented with the same images flipped horizontally. Though we effectively double our training size by flipping images, we do not always see improvements for doing so. The building and mean IU metrics are higher for the models fine-tuned on augmented datasets for 35 epochs, but are lower for models fine-tuned on augmented datasets for 100 epochs, as shown in Figure 4. At 35 epochs, models fine-tuned with the augmented dataset attained 45.29% and 72.62% building IU and mean IU, respectively, while models fine-tuned with the original dataset attained 38.65% and 69.30%. At 100 epochs,

models fine-tuned with the augmented dataset attained 46.65% and 73.30% building IU and mean IU, while models fine-tuned with the original dataset attained 47.19% and 73.58%. The first observation is likely the result of greater training size helping the FCNs to more rapidly fine-tune away from their previous ImageNet classification applications and towards the semantic segmentation of buildings, on a per epoch basis. The latter observation may be due to losses conferred by generalization, though these test result error metrics are also too close to be considered statistically significant. One possible explanation is that, because our WorldView-2 image has a off-nadir angle of 21° and was taken with asymmetrical sun azimuth and sun elevation of 111.5° and 75.3° , respectively, flipping the image will help the model to learn to detect buildings from different angles and with different lighting conditions. However, since we are only testing on images taken at constant angle and with constant lighting conditions, our specialized FCNs, trained for only these conditions, outperform more generalized versions. Further work is required to assess the potential benefits that data augmentation confers to model generalization ability.

3.4 Load Localization and Characterization

One of the problems with the pixel-based building extraction approaches discussed previously is that closely packed buildings are often clumped together into the same contiguous groups of pixels. This phenomenon can be observed when looking at the inferences presented in Fig. 3-10. Such clumping together makes it difficult to discern closely situated or connected buildings from single, larger buildings; this can adversely affect electricity infrastructure planning as individual buildings have different power loads and connection costs than single larger buildings. Two potential solutions exist in the literature, and we present additional methodologies here.

Varshney et al. describe a polygonization method aimed at fitting polygons onto individual buildings given a set of pixel-based inferences. This provides approximations on where distinct buildings exist, and provides a measure of their footprint sizes. Nevertheless, the authors note that the polygonization method is limited in its ability to distinguish between adjacent buildings (Varshney et al., 2015).

Yuan’s approach for describing buildings using a signed distance function from their boundaries in convolutional neural networks may be useful for this application. As mentioned in Section 3.1.4, Yuan’s output representation enables the definition of fine-grained labels for border boundaries (Yuan, 2016). Nevertheless, it remains to be seen whether this approach is efficacious for building extraction in rural areas of developing countries. The data set used by Yuan is comprised of 0.3 m resolution imagery for Washington D.C.; building rooftops in these images are much more clearly defined and easily distinguishable than they are for the 0.5-1.0 m resolution imagery that is more commonly available for the developing areas we are interested in. Furthermore, close inspection of Yuan’s inference results reveals that clumping is still prevalent when buildings are connected.

We introduce simpler approaches for distinguishing buildings from one another in satellite imagery and argue that they may be desirable given certain assumptions.

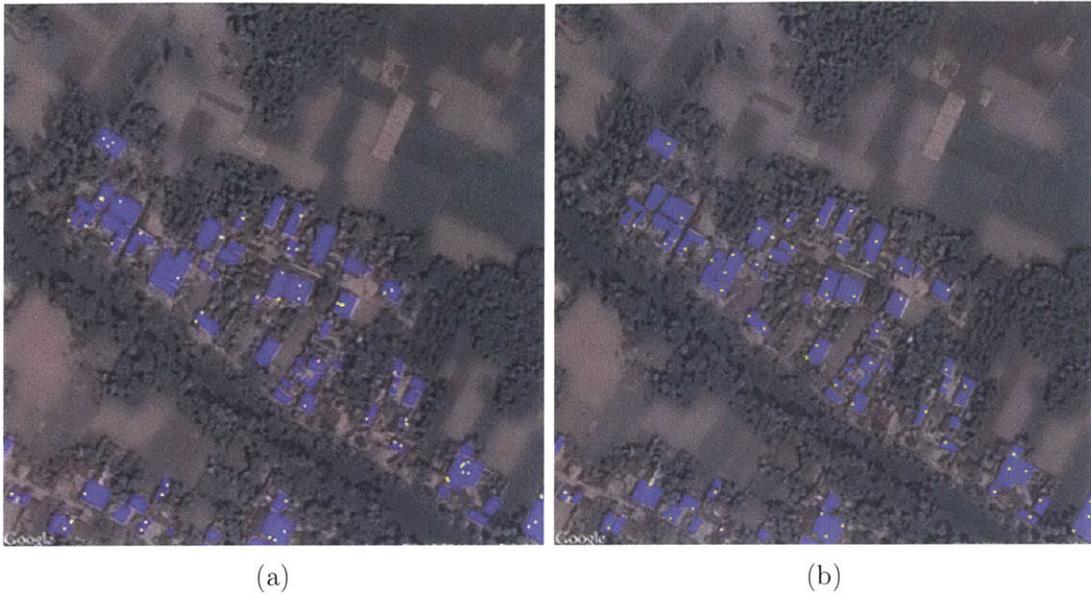


Figure 3-12: Load localization using variants of the same methodology: (a) uniformly sampling building pixels without replacement, and (b) uniformly sampling building pixels using exclusion regions

The first, naive approach, is exemplified by Fig. 3-12a, showing ground truth building footprints in blue and sampled connection points with yellow dots. Connection points are sampled from building pixels uniformly at random without replacement. The number of connection points per image is calculated using census data: the number of building pixels is divided by the number of buildings in a given administrative boundary to arrive at the number of pixels per building for the region. Subsequently, the number of building pixels in a given image tile is divided by the number of pixels per building and rounded to arrive at the number of building points that should be sampled. This approach approximates the number and dispersion of buildings on average, assuming all buildings are equally sized. Nevertheless, by looking at Fig. 3-12a, it is clear that this approximation is limited: individual buildings are not well represented by connection points and there does not seem to be even distributions of connection points throughout the building pixels.

To account for some of these deficiencies, we present a second related approach which performs uniform sampling of building pixels but also employs exclusion zones to promote the dispersal of building points. Algorithm 2 and Fig. 3-12b outline and demonstrate this methodology. Each time a point is sampled, we define a disk-shaped exclusion zone around it preventing the subsequent sampling of nearby points. The size of the exclusion disk is defined such that we can always sample the necessary number of connection points. Credit for this approach should be shared with Dr. Claudio Vergara, who assisted in the development of this feature. While Fig. 3-12b reflects visually improved results over Fig. 3-12a, it should be noted that this approach biases against sampling large contiguous regions of building pixels, as sampled pixels

within these regions have larger active exclusion zones than sampled pixels with fewer nearby building pixels. Stray building pixels output from the building extraction step are more likely to be sampled here than in the previous method. This qualitatively appears to be desired behavior; however, further research is necessary to validate this intuition.

Algorithm 2 Load localization with exclusion zones

Require: a binary *mask* constituting a matrix of building pixels

Require: n_b , the number of buildings to be sampled

- 1: Define n_p as the number of building pixels in *mask*
 - 2: Define *disk* as a disk structuring element s.t. its area is maximized while $\leq n_p/n_b$
 - 3: **for** $k = 0$ to n_b **do**
 - 4: Sample (i, j) uniformly at random from *mask* and save as a sampled point
 - 5: Define zeros matrix *exclude* with same dimensions as *mask*
 - 6: Set *exclude* index $(i, j) = 1$ and dilate *exclude* using *disk*
 - 7: For all indices where *exclude* is nonzero, set *mask* to zero values
-

While the two approaches we introduce provide easy ways to localize buildings from the building extractions obtained, they are limited by their necessary assumptions. They first assume that within a given administrative boundary, building extraction is equally sensitive to all buildings, and thus we can calculate the number of pixels per building for sampling as described above. This assumption may be acceptable given that image tiles within a given administrative boundary (from which we use census data to obtain building counts) are likely to be derived from the same original satellite image strips and thus have similar building extraction characteristics. When there are numerous diverse image strips within an administrative boundary, this assumption becomes less credible. In addition, when our building extraction methodologies are systemically more responsive to some types of buildings than others, this assumption poses difficulties. Indeed, we qualitatively observe that our FCN model is better able to detect buildings with metal roofs than thatched roofs. While this is problematic for accurate accounting, it may affect our electrification use-case less than other applications. We could expect the occupants of buildings with metal roofs to be more affluent than those with thatched ones and thus have higher latent power demand in practice; additional research is required to assess the extent to which these two biases are aligned.

A second assumption made by our load localization methodologies is that buildings in a given administrative boundary are of a single size and therefore should be considered equally sized electric loads. This assumption is more questionable, since it is clear from satellite imagery that buildings come in a range of sizes in our regions of interest. Nevertheless, it may be an acceptable approximation for electricity infrastructure planning if latent electric power demand scales roughly linearly with building rooftop area. For example, if a building has twice the area of an average building and has twice the latent power consumption, then modeling it as two separate buildings may not adversely affect subsequent electrification planning methodologies. This assumption is weaker, however, if we are comparing buildings with different uses and

with variable numbers of stories. Without a way of characterizing building height and economic uses, our load localization scripts are unable to account for these variables. Further research on using remote sensing techniques to infer these characteristics is required for improvement.

3.5 Large Scale Building and Load Localization

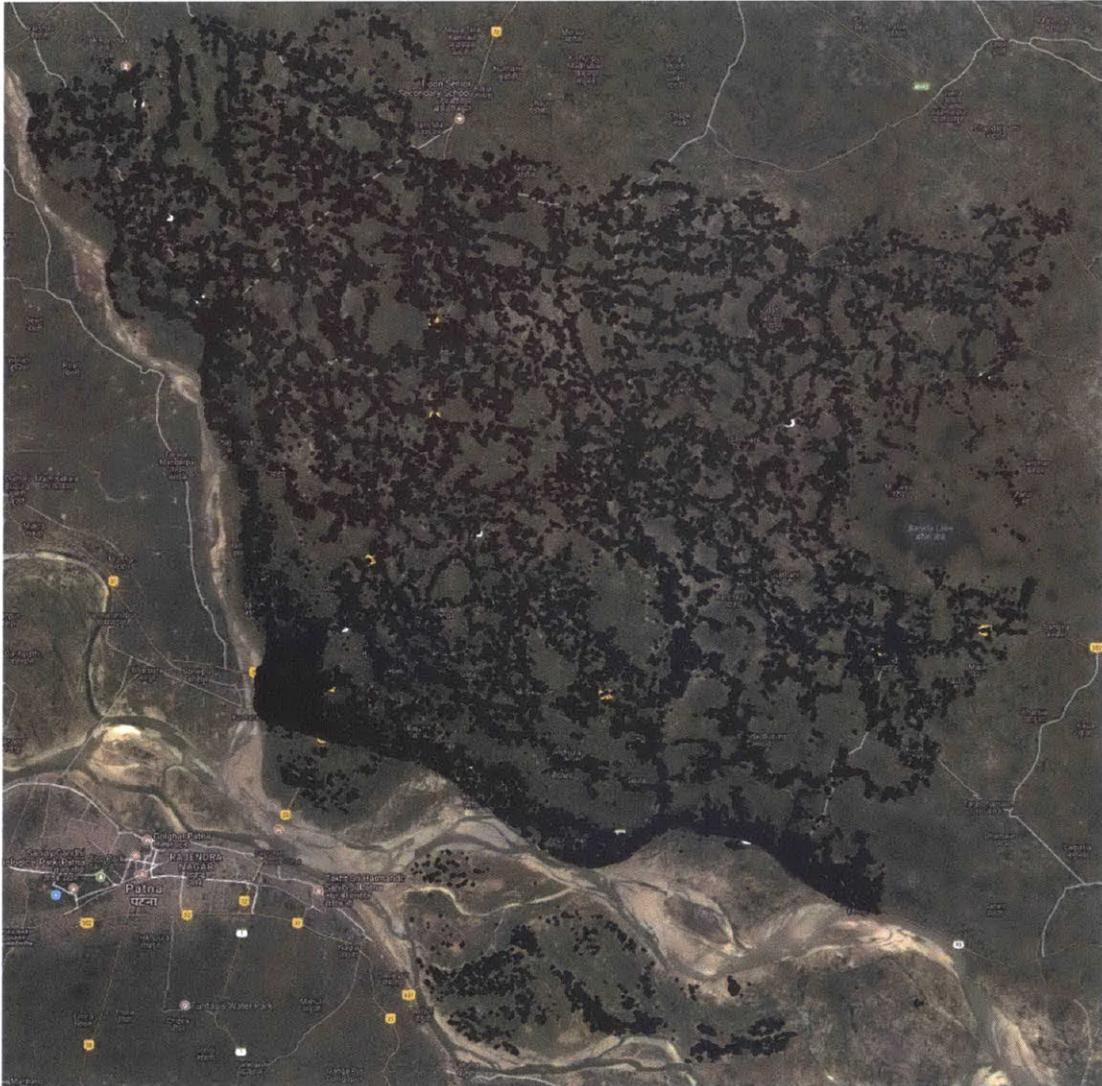


Figure 3-13: Large scale building identification for Vaishali, Bihar, India (>2000 km²)

With the building extraction and load localization methodologies discussed, we show that we are able to scale up to do large-scale building extraction for regions in India and Uganda. Fig. 3-13 and Section 4.2.2 show output from this process for the ~2000 km² Vaishali district in Bihar, India and the ~11,000 km² South Service

Territory for electric power in Uganda. In the South Service Territory, the building locations we produced missed large regions of buildings due to the nonavailability of satellite imagery from the Google Maps API. The classifier also had a moderate number of false positives and false negatives, and higher quality building locations were desired for planning. As a result, building locations were partially filled in and corrected by hand, using multiple different satellite data sources. We found that automatic building extraction can be a very valuable tool for energy infrastructure planning and lends itself well to hybrid automatic-manual approaches for geospatial planning.

3.6 Conclusions

This chapter introduces ongoing efforts to extract building footprints and localize building points from satellite imagery. ConvNets have emerged in recent years as the gold standard method for building extraction and a number of related computer vision tasks. Feedforward neural networks, ConvNets, and FCNs for semantic segmentation are introduced and experiments related to the procurement of training data are discussed. Notably, crowdsourcing building polygon annotations using Amazon Mechanical Turk is difficult and often leads to poor results. Instead, hiring “expert annotators” is shown to yield qualitatively higher quality training sets. Furthermore, simple experiments on building extraction using FCN models are reported. Using higher resolution and better quality images yields large improvements for building extraction with regards to error metrics. Data augmentation experiments are more mixed and motivate further analyses on generalization capabilities. Finally, simple algorithms are proposed for load localization and evidence for the scalability of deep learning systems for this application is presented.

Chapter 4

Electrification Status Estimation

Information regarding the current electrification status of buildings is imperative for energy access planning. This information allows planners to avoid planning for redundant infrastructure and ensure that their plans meet the full needs of the public. Electrification status information also enables the informed assessment of technology choices; planners can use it to determine the attractiveness of off-grid technologies relative to modes of grid extension. Extending the main grid is more economical near parts of the existing grid with high reliability. Microgrid and stand-alone system technologies are generally more attractive in areas far from the main grid.

Although distribution companies in developed parts of the world generally have a wealth of digitized infrastructure data, their counterparts in developing regions are consistently less informed. In 2016, we (the MIT Universal Access Lab) learned that our distribution company partners in India, Nigeria, and Uganda did not have or had incomplete structured information on its low-voltage distribution lines. Though data collection and digitization efforts have commenced in some of these regions since then, this data is still largely missing. Of the distribution companies we surveyed, only the Energy Development Corporation Limited (EDCL) of Rwanda, operating in a relatively small country, had a large fraction of low voltage data readily available.

This chapter will cover the development of machine learning methods for the estimation of electrification status using incomplete features related to electrification. Approaches span both discriminative and generative methods, including logistic regression, Gaussian processes, and Bayesian networks. Inference is performed using sparse electrification status surveys and commonly available data sources including nighttime lights imagery, building locations, census, survey data, and medium voltage-to-low voltage distribution transformer locations. To our best knowledge, this chapter represents the first reported studies of probabilistic approaches being employed for high resolution electrification status estimation. The overall probabilistic framework enables key aspects of planning and decision-making under uncertainty. We further discuss implications for decision-making under uncertainty in Chapter 6.

An important distinction needs to be made between electric power consumption and electrification status. A number of studies have focused on using some of these features to estimate electric power consumption in different regions of the world (Elvidge et al., 1997; Amaral et al., 2005; Chand et al., 2009; Townsend and Bruce, 2010; Letu

et al., 2010; Zhao et al., 2012; Shi et al., 2014). While electric power consumption is an important indicator for evaluating energy poverty in general, it has different planning implications from electrification status. Electric power consumption is only relevant once power systems infrastructure is present where people live. Its value is then affected by supply constraints and levels of consumer demand. On the other hand, building-level electrification status, a binary metric, measures whether or not power infrastructure exists, assuming that once it does, electric power of some reliability level will be available. When a building is electrified but has low electric power consumption due to reliability issues, a useful intervention may be to improve electric power generation capacities. When a building or region is non-electrified, grid extension or microgrid development may be necessary interventions. We find that the literature on performing electrification status estimation is much sparser than that on electric power consumption. While both measures are important for the provision of electric power, electrification status estimation is the focus of this section as it is more relevant to planning basic energy access and enabling the techno-economic planning tools that will be presented in Chapter 5.

Table 4.1: Methods for electrification status estimation

Method	Granularity	Regions Studied	Authors	Consumes Building Location	Consumes Nighttime Lights	Consumes Transformer/MV Grid Location	Consumes Electrification Survey	Consumes Aggregated Population	Consumes Aggregated Electrification
Nighttime lights Assumptions	~1km ²	World	Doll et al.		✓				
Logistic Regression	Village	Senegal; Mali; Vietnam	Min et al.		✓			✓	✓
Buffer areas	Building	Vaishali, India	Ellman			✓			
Artificial LV Network	Building	Vaishali, India	Cotterman		✓	✓			
Score-based	Building	SST, Uganda	Cotterman		✓	✓			
Logistic Regression	Building	SST, Uganda; Uganda Country	This thesis		✓	✓	✓		
Gaussian Process	Building	SST, Uganda; Uganda Country	This thesis				✓		
Hierarchical Beta	~1km ²	SST, Uganda; Kayonza, Rwanda	This thesis	✓	✓	✓	✓		✓

4.1 Background

In this section, background from the academic literature is presented relating to electrification status estimation, analyses using nighttime lights, the analysis of energy access near power transformers, and hierarchical spatial models. Electrification status estimation is the overall topic of this chapter and both nighttime lights and power transformer data is employed in the models presented. Finally, one of the contributions of this thesis is describing a novel hierarchical spatial model we develop for electrification status estimation. We explore the relevant literature to provide context for this work.

4.1.1 Electrification Status Estimation

With the exception of early precursors to the work presented in this thesis (i.e. models presented by Ellman and Cotterman), to our best knowledge, only two approaches for

electrification status estimation have previously been reported on in the literature, as illustrated in Table 4.1. Doll et al. estimate electrification status by making the assumption that anywhere with zero light intensity in Defense Meteorological Satellite Program - Operational Linescan System (DMSP-OLS) annual composite nighttime light images confers lack of electrification (Doll and Pachauri, 2010a). While Doll et al. are able to expand their analyses to very large regions with ease, their assumption is nevertheless questionable. Buildings that are located in areas that have positive nighttime light signals may be non-electrified and conversely, buildings that are located in areas with zero-valued nighttime light signals may be electrified. Fig. 4-9 depicts surveyed buildings and DMSP-OLS nighttime lights for the South Service Territory of Uganda, and displays both of these contradictory cases. Min et al. corroborate the inadequacy of Doll et al.'s assumption, finding that nighttime lights imagery most strongly reflects the presence of streetlights and is not on its own a strong indicator for household electricity use (Min et al., 2013a). As such, this methodology is unsuitable for the infrastructure planning activities we propose.

Min et al. consider energy access in Senegal and Mali at the village-level for the year 2011. They compare nighttime light output from the DMSP-OLS sensor against survey data representing 232 electrified and 899 unelectrified villages. Among other studies, the authors present a logistic regression model using population and monthly average light output to classify village electrification status. Though they produce a visualization conveying classification efficacy, they do not publish metrics for classification efficacy using this model. Furthermore, Min et al. do not provide the precise definitions they use for village electrification (Min et al., 2013a). As mentioned in Section 2.4, definitions of electrification are paramount to understanding electrification status and planning effective interventions. Two related considerations render the Min et al. methodology inadequate for detailed electrification planning: the aggregated village-level nature of electrification status presented and the use of a binary measure for village-level electrification status. Aggregated treatments of electrification status require disaggregation before they can be used to plan building-level power systems in heterogeneous environments. From what we have witnessed in our field visits and from what has been reported in the literature, "under grid" villages in developing countries are often highly heterogeneous (Lee et al., 2016). This makes the disaggregation of village-level metrics subject to inaccuracies. Binary metrics for village electrification status aggravate this consideration. Villages in the developing countries under consideration are usually not simply 0% or 100% electrified; even connected villages can have large populations without energy access. If village-level electrification status was instead interpreted as a continuous range between 0% and 100%, disaggregating village-level electrification measures for detailed infrastructure planning activities would be more feasible.

The electrification status estimation methodologies presented in this chapter have supplanted previous work performed by our group, including that by Ellman and Cotterman. Ellman makes the assumption that electrified buildings are purely contained within a set distance from medium voltage feeders. He defines "buffer" areas around the medium voltage feeders with extent proportional to the reported number of electrified buildings within an administrative region (Ellman, 2015). Ellman's approach

is visually questionable, however, as it is known that low voltage lines are heterogeneous and extend from transformers towards population centers. Furthermore, in many cases, buildings near low voltage lines and transformers may still be non-electrified. Cotterman illustrates two different methodologies that rely on more data. Cotterman first builds artificial low voltage networks around medium voltage-to-low voltage distribution transformers and defines electrified buildings to be buildings that exist in close proximity to these lines. This artificial network is designed using a set of low-voltage network templates; the templates are oriented at random angles around transformers to simulate a low voltage network (Cotterman, 2017). While this approach produces visually more dispersed results, its fabricated nature based on the random placement of predefined grid topologies undermines the spirit of detailed planning endeavors. Subsequently, Cotterman defines a score-based methodology using features including nighttime lights data, transformer locations, and high and medium voltage distribution lines. He then designates a share of buildings with the highest scores as electrified (Cotterman, 2017). While this approach encompasses more data and presents a forward model using expert-derived priors, it does not attempt to learn parameters or evaluate electrification status probabilistically.

4.1.2 Nighttime Lights

Nighttime lights imagery has been used to predict a number of economic indicators including GDP and electricity consumption, mostly in the developed world (Elvidge et al., 1997; Amaral et al., 2005; Doll et al., 2006; Chand et al., 2009; Townsend and Bruce, 2010; Letu et al., 2010; Zhao et al., 2012; Shi et al., 2014). It has also been used with daytime satellite imagery and convolutional neural networks to predict impoverished regions in the developing world (Jean et al., 2016). Researchers have found nighttime lights imagery valuable for a variety of remote sensing analyses due to their extensive coverage, ease of access, and quality. The DMSP-OLS dataset commonly used for these studies is described in detail in Section 4.2.1.

4.1.3 Distance to Transformers and Electrification

Lee et al. makes recommendations for electrification strategies supported by a novel data set of over 20,000 geo-referenced buildings across 150 Western Kenyan rural communities. They make the distinction between "off grid" and "under grid" households, and explain that policies focusing connection subsidies for "under grid" communities can significantly improve electrification outcomes. As part of their study, they plot grid connection rates by distance to transformer, building use, and wall quality (Lee et al., 2016). Their correlations demonstrate the utility of transformer distance information and motivate the incorporation of transformer distance features into the studies presented here.

4.1.4 Hierarchical Spatial Models and Hierarchical Models

Spatial modeling has its roots in the early 1900s and has since found myriad applications in a number of fields including economics, geology, ecology, atmospheric science, and epidemiology (Cressie, 1993). Spatial models using hierarchical Bayesian methods, however, are relatively newer and have grown in popularity along with MCMC methods in the last few decades. They enable the representation of and inference over rich and complex geospatial data sets (Banerjee et al., 2003). Among other applications, hierarchical Bayesian models have been successfully employed for multi-scale and multi-resolution inference.

Arab et al. describe that while hierarchical spatial models have matured, their inherent high dimensionality complicates the modeling process and can pose computational challenges. The authors stress the need for continued research on efficient computational methods (Arab et al., 2008). On this topic, Choi et al. show how pyramidal Gaussian graphical models can efficiently enforce statistical links between and within various scales (Choi and Willsky, 2007). Nevertheless, Gaussian graphical models are less suitable for non-Gaussian data, and Prates et al. and Yu et al. present methods using Gaussian copulas to extend these methodologies to represent multi-scale phenomena with non-Gaussian data (Prates et al., 2012; Yu et al., 2012).

The hierarchical beta model presented in Section 4.6 avoids the use of Gaussian copulas and presents efficient non-Gaussian hierarchical spatial models exploiting beta-binomial conjugacy. It takes advantage of architectural developments from hierarchical Bayesian models and employs aspects of empirical Bayes for the intelligent incorporation of multi-modal inputs. The hierarchical beta model also uses a Bayesian network architecture with multi-to-multi associations between groups of spatial data to exploit various levels of statistical dependency. Lin et al. present similar multi-to-multi associations, coupling mixture models for different data groups and sets of latent Dirichlet processes, and show how they enhance model flexibility and improve performance on document analysis and image modeling applications (Lin and Fisher, 2012).

4.2 Features Related to Electrification

In this section, various data sets are presented that are used in the case studies and electrification models covered in this chapter. As denoted in Table 4.1, analyses are presented for the South Service Territory for electric power in Uganda, the whole country of Uganda, and for the district of Kayonza, Rwanda.

4.2.1 Global Inputs

Nighttime Lights Data

We use nighttime lights data with global coverage in the form of annual composites from the DMSP-OLS for the year 2013. DMSP-OLS annual composites are



Figure 4-1: 2013 DMSP-OLS annual composite imagery

processed and made available by the National Oceanic and Atmospheric Administration’s (NOAA’s) National Geophysical Data Center ([National Geophysical Data Center, 2017](#)). The DMSP-OLS satellite orbits the earth with an altitude of 833 km, and images low-level radiation emissions at two wavebands: 0.4-1.1 μm and 10.5-12.6 μm , for detecting visible/near infrared and thermal infrared lights, respectively. The satellite typically makes a night-time pass between 20:00 and 21:30 each night, and processing is done to omit images with noise and to prevent signals from cloud cover, aurora glare, and other sources from influencing the composites. The resultant images of stable nighttime lights have pixels with brightness values ranging from 0 to 63. They are comprised of 30 arc second grids ($\sim 1 \text{ km}^2$), spanning -180 to 180 degrees of longitude and -65 to 75 degrees of latitude ([Doll and Pachauri, 2010b](#); [Min et al., 2013b](#); [National Geophysical Data Center, 2017](#)). Fig 4-1 shows the DMSP-OLS annual composite image for the year 2013, and Fig. 4-2 shows a view zoomed into the Uganda South Service Territory, which is discussed in the next subsection.

4.2.2 Inputs for the South Service Territory, Uganda

In this subsection, a variety of data sources are presented for the South Service Territory for electric power in Uganda (Uganda SST). The SST represents one of multiple electric power service territories that REA Uganda is producing master plans for. It constitutes 10,914 km^2 in the south of the country and includes part of Lake Victoria. Our partners at Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH and at the Rural Electrification Agency of Uganda (REA Uganda) provided transformer location, survey data, and census data for the region. These data sets represent the most diverse set of features analyzed in this thesis.

Building Location Data

Fully convolutional neural networks and building localization methodologies for the estimation of building locations within the Uganda SST are discussed in Section 3.5.

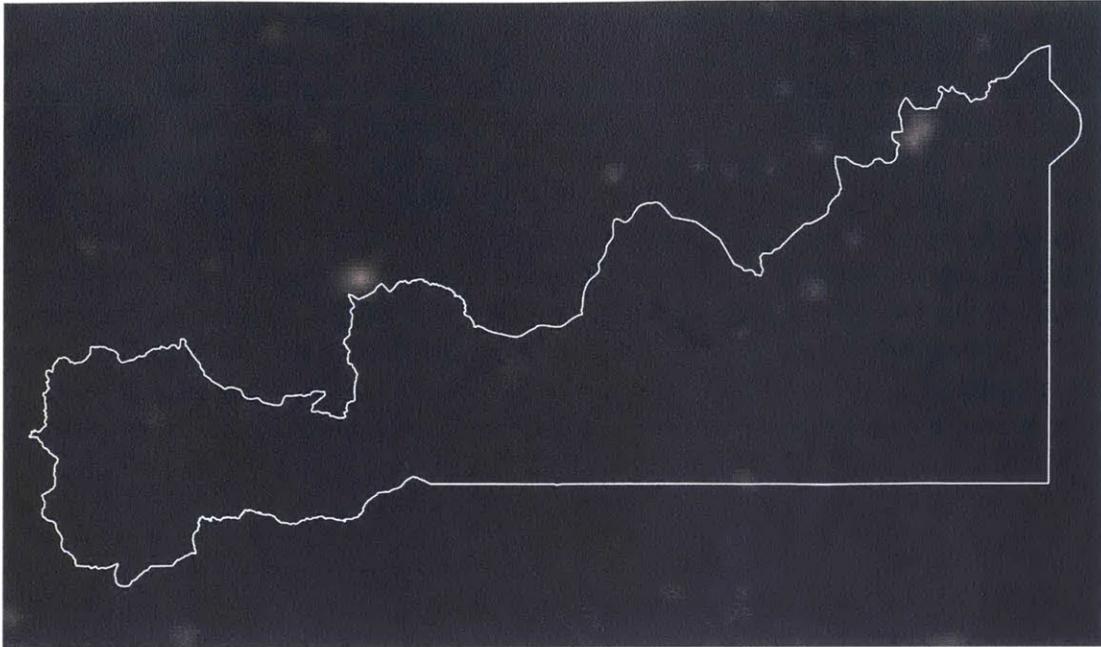


Figure 4-2: 2013 DMSP-OLS annual composite imagery for the Uganda SST

Our partners at GIZ and REA Uganda helped to provide Uganda-specific training data for this task, and we ran inference on imagery from the Google Maps API that was available as of July 2016. Note that imagery available through the Google Maps API on this date could potentially be several years old. Buildings locations are shown in Fig. 4-3. While most of the territory is covered by buildings, it is evident that there are several angular patches with missing buildings. This is due to the fact that Google Maps lacked building-level image strips for these regions, and as such, did not return corresponding image tiles. Analysis is performed on these areas recognizing that some regions reflect incomplete and outdated information.

Transformer Locations and Distances, 2016

While REA Uganda provided high voltage distribution line, medium voltage distribution line, and medium-to-low voltage transformer data for the Uganda SST, only transformer location data is used. This is because we assume the topology of the low voltage network is conditionally independent of the medium and high voltage networks given the transformer locations. Low voltage lines necessarily connect to transformers and they are distributed according to the locations of consumers, not to be around medium or high voltage lines. Transformer locations are shown as blue dots in Fig. 4-4, along with an outline of the Uganda SST border and sub-county borders. REA Uganda reported that the transformer data set was mostly complete as of late 2016; however, transformer data is likely still missing in some areas.

As an additional feature, transformer distance maps are computed for every point along a 30 arc second grid covering the Uganda SST. This distance map uses a com-



Figure 4-3: Buildings extracted and localized using FCN models, as described in Section 3.5.

mon k-nearest neighbors algorithm (k-NN) made available in the scikit-learn machine learning library for Python (Pedregosa et al., 2011).

Sparse Survey Data, 2016

GIZ conducted an electrification survey of the Uganda SST in 2016 with our guidance. GIZ surveyed both urban and rural areas of Uganda on motorbike and targeted predefined clusters of about 10 buildings each that were chosen semi-randomly. Clusters were defined such that they were situated around buildings and they were distributed so that they covered the populated regions of the Uganda SST fairly evenly. Enumerators were instructed to visually inspect buildings in these clusters and determine whether they had overhead power connections. If the enumerator suspected that underground connections were present, he or she was instructed to inquire about the building's electrification status to validate it. As a result of this partially subjective process, it is acknowledged that survey error may be biased in favor of false negatives due to the presence of underground lines in urban areas; nevertheless, this problem is considered infrequent. Enumerators used Android tablets with GPS capabilities and a survey form developed on the Open Data Kit platform (Hartung et al., 2010). Overall, the data set contains 209 electrified buildings and 263 non-electrified buildings. It is visualized in Fig. 4-5.

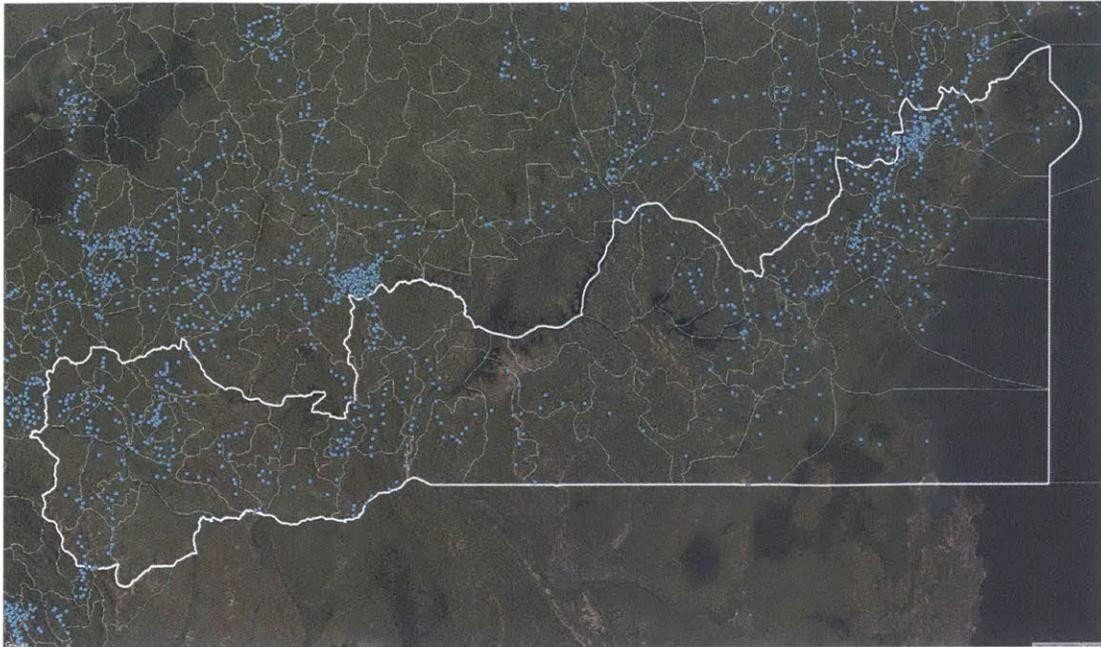


Figure 4-4: Uganda SST transformers are depicted with blue circles, the Uganda SST administrative boundary is shown with a bold white border, and sub-country boundaries are shown with a faint white border.

4.2.3 Inputs for Uganda at the Country-Level

Features for the 241,038 km² country of Uganda were available including aggregate census data, transformer location information, and a different electrification survey data set from 2012. This data is used in Section 4.4 and 4.5 to evaluate electrification status using logistic regression and Gaussian processes, respectively. Aggregate census data is also employed for Uganda SST-specific analysis in Section 4.6 covering the hierarchical beta model.

Aggregate Census Data, 2014

The sub-county borders shown in Fig. 4-4 and Fig. 4-5 are useful because they allow the mapping of census data to geographic coordinates. Census data containing sub-county-level aggregate statistics is provided by the Uganda Bureau of Statistics (UBOS) for the year 2014 (Uganda Bureau of Statistics, 2014). While the census provides information about the spatial distribution of a number of demographic factors, we focus on two metrics: one concerning the number of households with electric lighting, and the other being the number of households in the sub-county. Sub-county electrification rates are approximated by computing the fraction of households that use electric lighting in a given sub-county. Our contacts at GIZ and REA Uganda have verified that this formulation is credible, as the first appliances purchased by electrified households are nearly always electric lights.

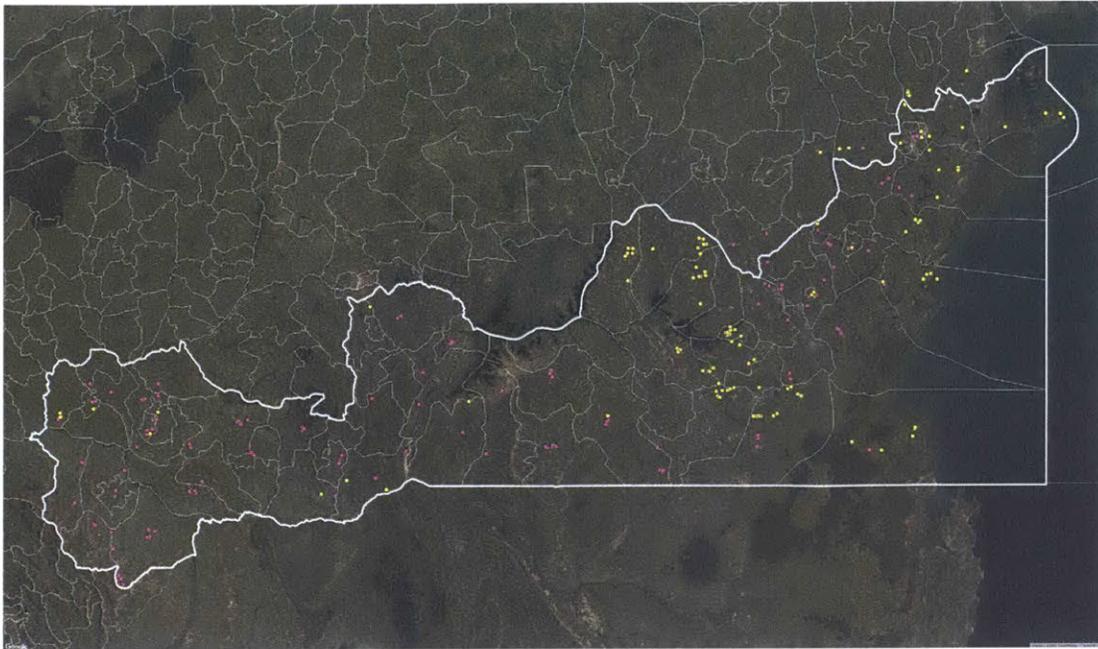


Figure 4-5: Electrified buildings for the Uganda SST from the 2016 GIZ electrification survey are depicted with yellow circles. Non-electrified buildings are represented by magenta circles.

Transformer Locations, 2011

Transformer locations were also provided by REA Uganda, but from the year 2011. While this data set has a much larger extent than the Uganda SST transformer data set, our contacts at REA Uganda described that it could have an even greater share of missing transformers than that for the SST. Fig. 4-6 shows medium-to-low voltage transformers as black circles.

Sparse Survey Data, 2012

Sparse electrification survey data for Uganda at the country-level in 2012 is provided by UBOS and the Uganda Ministry of Energy and Mineral Development. The survey covers 111 districts in Uganda and was intended to assess electrification levels within the country, provide indicators for socioeconomic sectors, and to ultimately support infrastructure planning (Ministry of Energy and Mineral Development and Uganda Bureau of Statistics, 2014). While the presence of multiple energy sources were assessed for households, businesses, educational facilities, and health buildings, we focus only on the presence or absence of electricity for households surveyed. Other facility types were excluded to avoid biasing representations of electrification probability, as different facility types had different sampling rates and are likely to have different electrification characteristics. Electrified and non-electrified households are shown in Fig. 4-6 as yellow and purple circles, respectively.

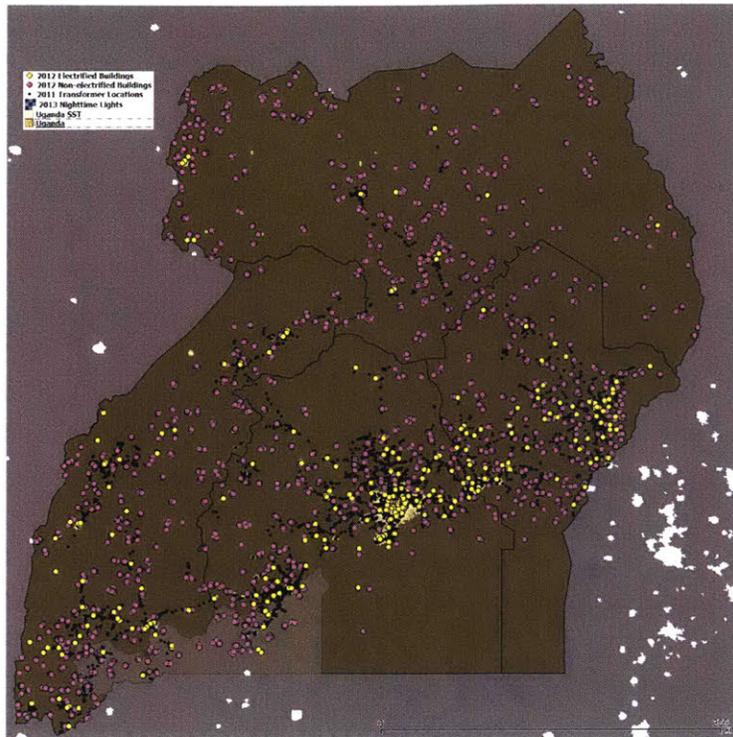


Figure 4-6: Electrified buildings from the Uganda country-level 2012 ERT survey are shown as yellow circles, non-electrified buildings from this same survey are depicted as purple circles, transformers are the black circles, and nighttime lights comprise the transparent layer. The Uganda SST administrative boundary is also shown in the south of the country to provide a sense of scale.

4.2.4 Inputs for for Kayonza, Rwanda

The last case study presented in this thesis pertains to the 1,937 km² Kayonza district in the Eastern Province of Rwanda. Though this region is small relative to the Uganda SST or Uganda at the country-level, it is valuable because we were able to obtain ground-truth low voltage network information from our partners at EDCL.

Building Location Data

Rwandan building location data was provided by a study conducted by French consulting firm Sofreco. The study, entitled "EWSA, Electricity Access Roll Out Program" was published in 2013 and is accompanied by a detailed geographic information system (GIS) database of consumer and infrastructure data relevant to energy access planning. Sofreco describes that they produced the consumer data by acquiring orthophotos for the whole country of Rwanda and then manually identifying 1,704,749 buildings across the country. While the study was released in 2013, the orthophotos used were taken 2-3 years earlier (Sofreco, 2013). Sofreco's building data set for Kayonza is shown in Fig 4-7 as black points.

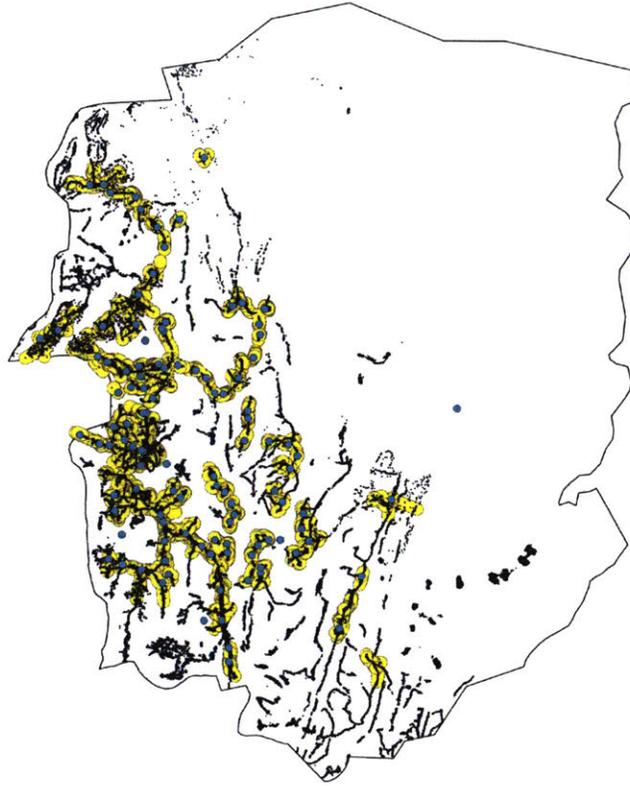


Figure 4-7: Low voltage grid, transformer locations, and buildings for Kayonza, Rwanda. The low voltage grid is highlighted with thick yellow lines, the transformers are shown with cyan dots, and buildings are represented by black points.

Transformer Locations, 2017

Updated transformer locations in Kayonza were obtained from EDCL and are considered complete as of July 2017. They are plotted as blue circles in Fig. 4-7.

Low Voltage Network

EDCL also provided low voltage network information which was estimated to be 90% complete as of 2016. It is plotted using large yellow lines in Fig. 4-7 to enhance its visibility.

Aggregate Census Data and Derivatives for Electrified Buildings

Aggregate census data on electrification rates for Kayonza is provided by the National Institute of Statistics in Rwanda. Kayonza was reported to have a 16.10% electrification rate in 2012 (National Institute of Statistics Rwanda, 2014). We use this statistic in concert with the low voltage distribution data and population data to provide derivative measures of electrified buildings in Kayonza. The low voltage network data for Kayonza is converted to a set of points representing the vertices of

the lines. A k-NN algorithm is then used to determine the distance between each of the buildings to the low voltage grid. Finally, the nearest 16.10% of buildings are designated as electrified, in accordance with the district electrification rate. We argue that this is a reasonable approach to determining ground truth electrification status. The assumption that only buildings nearby to the low voltage grid are electrified is sound because the off-grid electrification rate is extremely low and buildings cannot possibly be grid connected without connection to low voltage lines. Furthermore, the assumption that the low voltage line can be represented by a set of vertex points is reasonable since the representation of the low voltage grid is very granular and includes line segment lengths on the order of building lengths. This derivative set of electrified buildings is by far the best and most complete that we have acquired for a developing country. It should be noted, however, that some representation error is still likely given the fact that only 90% of the low voltage network is included in the EDCL data set.

4.3 Electrification Status Models and Error Rates

The next three sections describe model candidates for high granularity electrification status estimation. Common logistic regression and Gaussian process models are presented for building-level classification and a novel hierarchical beta model is presented for inference over 30 arc second grids (~ 1 km). For building-level classification, common error metrics including accuracy, precision, recall, and F_1 score are used to evaluate model performance with regards to survey data. Accuracy, precision, recall, and the F_1 score are given by 4.1 through 4.4 where P , N , stand for number of positive and negative classifications, respectively. Likewise, TP , TN , FP , and FN , stand for the number of true positives, true negatives, false positives, and false negative classifications.

$$\frac{TP + TN}{P + N} \tag{4.1}$$

Accuracy

$$\frac{TP}{TP + FP} \tag{4.2}$$

Precision

$$\frac{TP}{TP + FN} \tag{4.3}$$

Recall

$$\frac{2TP}{2TP + FP + FN} \tag{4.4}$$

F_1 score

Because the hierarchical beta model has an aggregate representation of electrification status, we conversely consider data likelihood scores for determining performance.

The data sets employed in the next three sections for Uganda, the Uganda SST, and Kayonza, are discussed in 4.2 and reflect features thought to be correlated with electrification status. They reflect measurements often taken at different times and in different years. Though techniques for spatio-temporal modeling exist, we make the assumption that data sets employed within any given study are contemporaneous and treat changes due to temporal discrepancies as measurement noise. This simplifies modeling and we assume the associated error is small if that the state of electrification has not changed significantly over the time spans considered for our region of interest. We explore how this assumption affects different model types, and discuss how models that are more robust to incomplete data sets are likely also more robust to misalignment in measurement dates.

4.4 A Logistic Regression Model

In this section, a logistic regression approach to building-level electrification status estimation is presented using nighttime lights and transformer distance values. While the relative simplicity of the model and features represented confer high model bias, the model still proves to be insightful and useful for benchmarking. It also motivates more complicated analyses.

4.4.1 Background

The Bernoulli Distribution

The Bernoulli distribution is a probability mass function (pmf) concerning a single binary random variable, $x \in \{0, 1\}$, which is commonly compared to a coin toss. $x = 1$ represents ‘heads’ or ‘success’ and $x = 0$ represents ‘tails’ or ‘failure.’ The probability that $x = 1$ is represented by

$$\Pr\{x = 1|\theta\} = \theta \tag{4.5}$$

where $\theta \in [0, 1]$. It follows that $\Pr\{x = 0|\theta\} = 1 - \theta$. The probability distribution for x is given below.

$$p_x(x) = \text{Bernoulli}(x; \theta) \triangleq \theta^x(1 - \theta)^{1-x} \tag{4.6}$$

This distribution has mean and variance given by 4.7 and 4.8, respectively.

$$\mathbb{E}[x] = \theta \tag{4.7}$$

$$\text{var}(x) = \theta(1 - \theta) \tag{4.8}$$

Logistic regression

Logistic regression is an example of a discriminative classifier because it fits models of the form $p(y|\mathbf{x})$, directly mapping inputs \mathbf{x} to outputs y . It specifically corresponds

to the following binary classification model:

$$p_{y|\mathbf{x},\mathbf{w}}(y|\mathbf{x},\mathbf{w}) = \text{Bernoulli}(y|\text{sigmoid}(\mathbf{w}^\top \mathbf{x})) \quad (4.9)$$

where $\text{sigmoid}(a)$ is the logit sigmoid function, defined as

$$\text{sigmoid}(a) \triangleq \frac{1}{1 + \exp(-a)}. \quad (4.10)$$

Maximum likelihood estimation is used to determine the parameters \mathbf{w} . The likelihood function is represented in 4.11, and we derive the negative log likelihood as 4.12.

$$p(\mathcal{D}|\mathbf{w}) = \prod_{i=1}^N \left[\text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i)^{y_i} (1 - \text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i))^{(1-y_i)} \right] \quad (4.11)$$

$$\text{NLL}(\mathbf{w}) = - \sum_{i=1}^N \left[y_i \log \text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i) + (1 - y_i) \log(1 - \text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i)) \right] \quad (4.12)$$

To minimize the negative log likelihood, we compute the gradient as 4.13 and apply an iterative optimization method such as gradient descent or iterative reweighted least squares. Because the error functions are convex, the optimization methods discussed can find global optimum values.

$$\frac{d}{d\mathbf{w}} \text{NLL}(\mathbf{w}) = \sum_i (\text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i) - y_i) \mathbf{x}_i \quad (4.13)$$

While global optimum values are attained, it is often still desirable to use regularization techniques to improve model generalization. l_2 regularization is commonly used for this purpose. Under l_2 regularization, we modify the objective and gradient functions as shown in 4.14 and 4.15, respectively.

$$f(\mathbf{w}) = \text{NLL}(\mathbf{w}) + \lambda \mathbf{w}^\top \mathbf{w} \quad (4.14)$$

$$\frac{d}{d\mathbf{w}} f(\mathbf{w}) = \sum_i (\text{sigmoid}(\mathbf{w}^\top \mathbf{x}_i) - y_i) \mathbf{x}_i + 2\lambda \mathbf{w} \quad (4.15)$$

4.4.2 Uganda Case Study

A logistic regression implementation for electrification status estimation is presented for the country of Uganda based on 2013 DMSP nighttime lights data as described in Section 4.2.1, transformer data from 2011 as detailed in Section 4.2.3, and sparse 2012 electrification survey data from Section 4.2.3.

While intensity values from the nighttime lights data set are used at face-value, distance features from the transformer data set are derived. Specifically, first-nearest

neighbor distances for each survey building are computed from the set of transformers using a k-NN implementation. The survey is broken up into a training set with a random 70% of the original buildings and a test set with the remainder. A validation set is not employed, because as we will describe, hyperparameter optimization is not attempted. With nighttime light intensity and transformer distance as features, l_2 regularized logistic regression is performed to classify building electrification status using stochastic gradient descent (SGD).

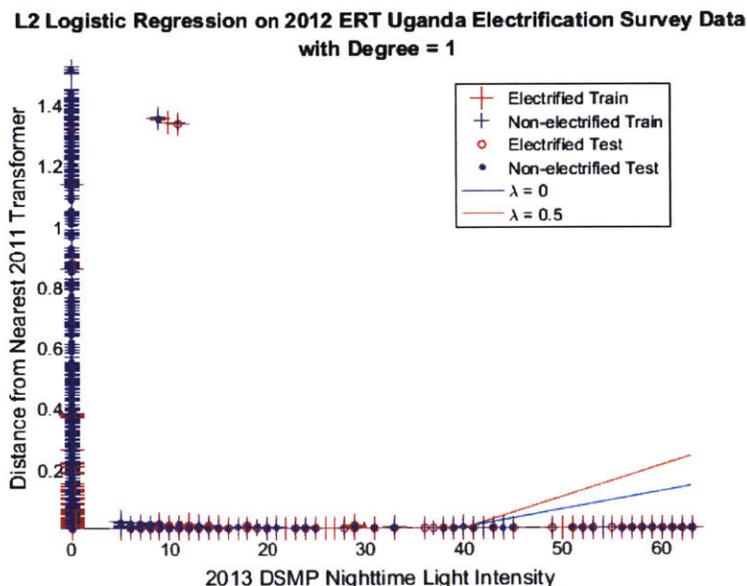


Figure 4-8: Logistic regression separating lines for the Uganda case study

Logistic regression and l_2 regularized separating lines are plotted in Fig. 4-8, showing distance from nearest 2011 transformer vs. 2013 DSMP nighttime light intensity. In addition, it is useful to refer back to Fig. 4-6, showing an illustration of the raw data sets used for the case study. Because of the nature of the correlation between the two dimensions, we find most buildings with nonzero transformer distance to have zero light intensity and most buildings with positive light intensity to be very close to a transformer. This correlation makes it such that we are essentially learning a one-dimensional classifier; in this case, we are classifying exclusively on nighttime light intensity. Buildings in regions with nighttime light intensity greater than 40 are classified as electrified, and buildings with nighttime light intensities otherwise are not. It can be qualitatively assessed that there is high overlap between 2013 nighttime lights data and 2011 distribution transformer locations for the whole country. This makes sense intuitively, since transformers are required to power streetlights and other nighttime light-inducing infrastructure. These trends are still evident despite the incomplete nature of the transformer data set, in addition to the discrepancy in the data set time periods.

Having a positive l_2 regularization parameter λ changes the characteristics of the two-dimensional decision boundary; however, it is empirically shown to have no net

effect on classification accuracies with this data representation due to the separable nature of the data. Because of the limited utility of regularization, hyperparameter optimization and performance metric calculation for regularized logistic regression are not pursued.

Error metrics for logistic regression are presented in Table 4.2, along with metrics for an all-negative classifier for comparison. It should be noted that logistic regression confers accuracy improvements over the all-negative classifier and that some degree of overfitting is exemplified when comparing accuracy and F_1 score values for training and test sets. The fact that we achieve low recall and F_1 scores overall is due to the fact that the two features we use do a poor job separating electrified from non-electrified buildings. This can be observed in the plots as well. While most electrified buildings are shown to be situated near transformers and lighted areas, there are still significantly many non-electrified buildings with these same characteristics.

Table 4.2: Logistic regression error metrics for the 2012 Uganda case study

Method	Data Set	Acc.	Prec.	Rec.	F_1 Score	Obj. Val.
All Neg.	Train	87.65	N/A	N/A	N/A	N/A
LR (Unreg.)	Train	89.21	0.63	0.30	0.41	1778.60
All Neg.	Test	87.64	N/A	N/A	N/A	N/A
LR (Unreg.)	Test	88.99	0.63	0.26	0.37	796.70

4.4.3 South Service Territory, Uganda Case Study

Logistic regression is run for the Uganda SST in a similar fashion as in Section 4.4.2; however, the study area is now confined to the SST and more recent data sets for transformer locations and surveys are used. As before, 2013 DMSP nighttime lights imagery is used and are described in Section 4.2.1. Instead, however, transformer data for the SST from 2016 is employed as detailed in Section 4.2.2, and sparse 2016 electrification survey data is used as presented in Section 4.2.2.

Data for the Uganda SST case study is plotted in 4-9. Unlike the Uganda case study, there does not appear to be compelling and positive correlation between electrification status, transformer vicinity, and nighttime light intensity. There are many high-light intensity areas with only non-electrified buildings surveyed, there are electrified buildings without transformers nearby, and there are non-electrified buildings surrounded by transformers.

A separating line from unregularized logistic regression is depicted in Fig. 4-10. In disagreement with the Uganda country-level case study, the logistic regression model in the Uganda SST case study chose nearest transformer distance as the best dimension for one-dimensional classification. Counterintuitively, the model specifies that buildings far away from transformers should be classified as electrified. Table 4.3 summarizes logistic regression model results compared to the all-negative classifier and corroborates this observation. Logistic regression model results are shown to improve on the all-negative classifier for both training and test cases. While the classifier

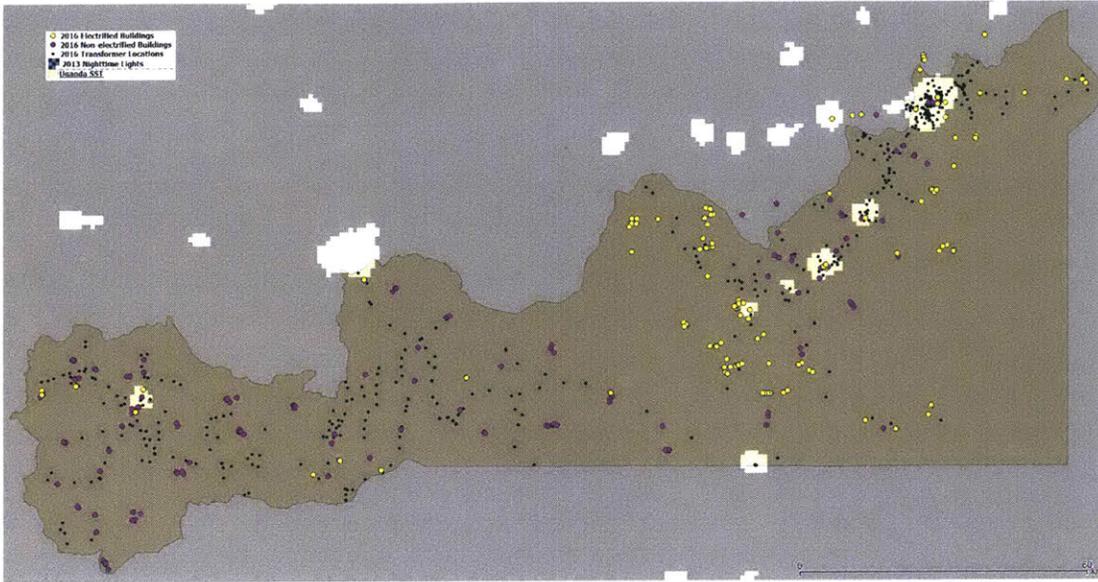


Figure 4-9: Uganda SST survey locations, nightlights, and transformers. Electrified buildings are shown in yellow, non-electrified buildings in purple, and transformer locations in black. Nighttime lights are overlaid, with brighter pixels representing higher intensity light.

seems contradictory from the viewpoint of commonsense, it is able to discern weak data set consistent characteristics: buildings farthest from transformers have a higher probability of being electrified. This is likely the result of missing transformers in the REA data set. Nevertheless, improvements from this feature are very weak, and the model only achieves low recall and F_1 metrics.

We note that error metrics for the Uganda country-level case study are much more favorable than those shown in this study for the Uganda SST. This could be attributed to the fact that the Uganda country-level case study is significantly larger in scale than that for the Uganda SST. With a larger case study, border effects are smaller, measurement noise is more distributed, and more data is available for training. Furthermore, data sets used in the Uganda country-level case study are collected closer together in time than those used for the Uganda SST study. Since electrification is a dynamic process, time discrepancies between features can result in noise and systematic error. The apparent weaknesses of this approach with the features presented motivate the search for other methods which may be able to more easily incorporate a richer set of features for electrification status estimation.

4.5 A Gaussian Process Model

In this section, a method for building-level electrification status estimation is presented based on Gaussian processes (GPs) for Uganda and the Uganda SST based exclusively on sparse survey data. While the logistic regression approach presented

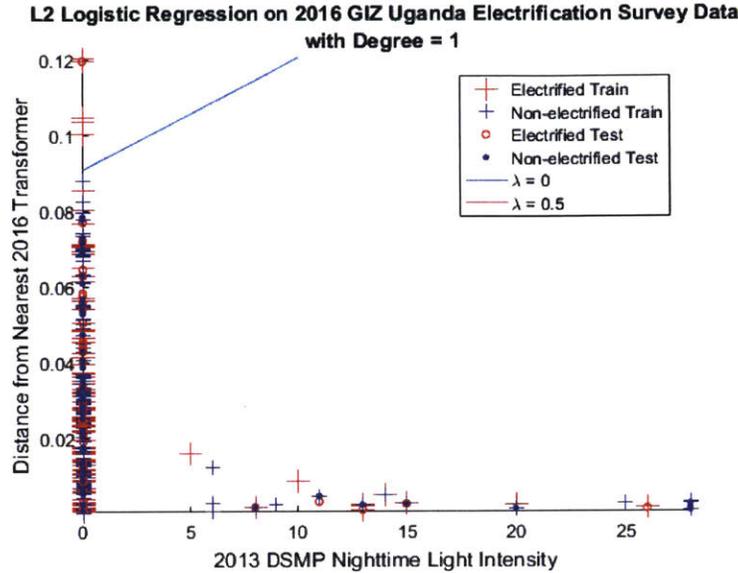


Figure 4-10: Logistic regression separating line for the Uganda SST case study

Table 4.3: Logistic regression error metrics for the 2016 Uganda SST case study

Method	Data Set	Acc.	Prec.	Rec.	F_1 Score	Obj. Val.
All Neg.	Train	55.76	N/A	N/A	N/A	N/A
LR (Unreg.)	Train	57.58	1.00	0.04	0.08	226.10
All Neg.	Test	55.63	N/A	N/A	N/A	N/A
LR (Unreg.)	Test	57.04	1.00	0.03	0.06	97.20

in Section 4.4 was able to easily incorporate nighttime lights and raw transformer distance features, Gaussian process classification (GPC) is naturally used to model spatial correlations in electrification status. While we find that the GPC approach has weaknesses with regards to its flexibility and ability to generalize to areas without survey data, it elucidates the importance of spatial correlation for electrification status estimation and motivates the pursuit of different and more flexible spatial models for this application.

4.5.1 Background

The basic theory for GPs goes back to at least the 1940s, and since then, GPs have found applications in different areas including geostatistics, meteorology, and circuit design. GPs are non-parametric methods that represent a generalization of the Gaussian probability distribution applied to stochastic processes. They are used for both regression and classification tasks in supervised learning and their consistency and computational tractability are often seen as strengths (Rasmussen and Williams, 2006).

We use GPs applied to the binary classification task of electrification status esti-

mation in this section. For this application, we are given geospatial data points \mathbf{x}_i with class labels $y_i = \{-1, +1\}$, and we try to predict class membership probabilities for a query point \mathbf{x}_* . GPCs achieve this by modeling a latent function f . Similar to logistic regression, a sigmoid function is used to map a function, in this case f , to the unit interval. As a result, the class membership probability can be given as $p(y = +1|\mathbf{x}) = \text{sigmoid}(f(\mathbf{x}))$. Because class membership must normalize and the sigmoid function satisfies the point symmetry condition, $\text{sigmoid}(t) = 1 - \text{sigmoid}(t)$, and we can write $p(y|\mathbf{x}) = \text{sigmoid}(y \cdot f(\mathbf{x}))$.

In accordance with the notation used in (Nickisch and Rasmussen, 2008), we define $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ as a matrix of training points, $\mathbf{y} = [y_1, \dots, y_n]^\top$ as a vector of target values, and $\mathbf{f} = [f_1, \dots, f_n]^\top$ as a vector of latent function values where $f_i = f(\mathbf{x}_i)$. We also define $\mathcal{D} = (\mathbf{X}, \mathbf{y})$ and refer to quantities with asterisks as query or test points. The factorial likelihood for the Bernoulli distributed data points is defined by 4.16.

$$p(\mathbf{y}|f) = p(\mathbf{y}|\mathbf{f}) = \prod_{i=1}^N p(y_i|f_i) = \prod_{i=1}^N \text{sigmoid}(y_i f_i). \quad (4.16)$$

Rasmussen et al. describe GPs as "a collection of random variables, any finite number of which have a joint Gaussian distribution," and show how GPs can be fully specified by a mean function $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ and a covariance function $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))^\top]$. As such, GPs can be written as $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ (Rasmussen and Williams, 2006) with parameters $\boldsymbol{\theta}$. A random variable $f(\mathbf{x})$ corresponds to every \mathbf{x} so that the joint distribution $p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{f}|\mathbf{m}, \mathbf{K})$ where $\mathbf{m} = [m(\mathbf{x}_1), \dots, m(\mathbf{x}_N)]$ and \mathbf{K} is a matrix with elements $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. Without loss of generality we set $\mathbf{m} = \mathbf{0}$.

The general equations specifying GPs for binary classification can be summarized by a few key equations, assuming zero mean. They are outlined below and described in greater detail in (Rasmussen and Williams, 2006) and (Nickisch and Rasmussen, 2008). By applying Bayes' rule, the posterior distribution over \mathbf{f} is given by 4.17.

$$p(\mathbf{f}|\mathbf{y}, \mathbf{X}, \boldsymbol{\theta}) = \frac{p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta})}{\int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta})d\mathbf{f}} = \frac{\mathcal{N}(\mathbf{f}|\mathbf{0}, \mathbf{K})}{p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta})} \prod_{i=1}^N \text{sigmoid}(y_i f_i). \quad (4.17)$$

For a set of N_* test points $\mathbf{X}_* = [\mathbf{x}_{*1}, \dots, \mathbf{x}_{*N_*}]$ with corresponding latent function values in \mathbf{f}_* we can write the joint and conditional distributions in 4.18 and 4.19.

$$p(\mathbf{f}_*, \mathbf{f}|\mathbf{X}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta}) = \mathcal{N}\left(\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} \middle| \mathbf{0}, \begin{bmatrix} \mathbf{K} & \mathbf{K}_* \\ \mathbf{K}_*^\top & \mathbf{K}_{**} \end{bmatrix}\right) \quad (4.18)$$

$$p(\mathbf{f}_*|\mathbf{f}, \mathbf{X}_*, \mathbf{X}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{f}_*|\mathbf{K}_*^\top \mathbf{K}^{-1} \mathbf{f}, \mathbf{K}_{**} - \mathbf{K}_*^\top \mathbf{K}^{-1} \mathbf{K}_*) \quad (4.19)$$

We write the joint posterior, $p(\mathbf{f}_*, \mathbf{f}|\mathbf{X}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta})$, as the product of the conditional

prior and the posterior, and marginalize it over \mathbf{f} to get 4.20.

$$p(\mathbf{f}_*|\mathbf{X}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta}) = \int p(\mathbf{f}_*, \mathbf{f}|\mathbf{X}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta})d\mathbf{f} = \int p(\mathbf{f}_*|\mathbf{f}, \mathbf{X}_*, \mathbf{X}, \boldsymbol{\theta})p(\mathbf{f}|\mathbf{y}, \mathbf{X}, \boldsymbol{\theta})d\mathbf{f} \quad (4.20)$$

Finally, the predictive class probability for a single point \mathbf{x}_* , $p(y_*|\mathbf{x}_*, \mathbf{X}, \boldsymbol{\theta})$, is obtained using 4.21.

$$p(y_*|\mathbf{x}_*, \mathbf{X}, \boldsymbol{\theta}) = \int p(y_*|f_*)p(f_*|\mathbf{x}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta})df_* = \int \text{sigmoid}(y_*f_*)p(f_*|\mathbf{x}_*, \mathbf{y}, \mathbf{X}, \boldsymbol{\theta})df_* \quad (4.21)$$

The integral from 4.21 is analytically tractable using the probit sigmoid function defined in 4.22, while it must be approximated when using the logit sigmoid function, as covered in 4.10.

$$\text{sigmoid}_{\text{probit}}(t) \triangleq \int_{-\infty}^t \mathcal{N}(\tau|0, 1)d\tau \quad (4.22)$$

4.5.2 Methods

Probabilistic binary classification using GPs is performed using the GPML MATLAB implementation. Model attributes are chosen based on recommendations from the literature and characteristics specific to the case study being explored (Rasmussen and Nickisch, 2016b). Characteristics of the implementation are provided below regarding the likelihood, mean, and covariance functions, in addition to covariance function and Laplace approximations.

- The likelihood function is used to define class membership probability and is defined to be $p(y|\mathbf{x}) = \text{sigmoid}(y \cdot f(\mathbf{x}))$. For the model presented, the sigmoid function is chosen to be probit error function, which confers a stronger penalty for misclassifications relative to the logit-based sigmoid function (Nickisch and Rasmussen, 2008).
- The mean function is chosen to be a constant to allow for greater model flexibility during training relative to the zero mean assumption. As such, $m(\mathbf{x}) = c$ where $c \in \mathbb{R}$ and is treated as a hyperparameter.
- A covariance function $k_\psi : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ with hyperparameters ψ is defined over the domain \mathcal{X}^2 . The function computes the covariance $k(\mathbf{x}, \mathbf{z})$ of \mathbf{f} between inputs \mathbf{x} and \mathbf{z} . A squared exponential covariance function was chosen with automatic relevance determination: $k(\mathbf{x}, \mathbf{z}) = \sigma_f^2 \exp(-1/2(\mathbf{x} - \mathbf{z})^\top \Lambda^{-2}(\mathbf{x} - \mathbf{z}))$ specifying hyperparameters ψ , λ_1 , λ_2 , and σ_f (Rasmussen and Nickisch, 2016b). Automatic relevance determination (ARD) uses separate length-scales for each dimension of \mathbf{x} and helps to determine how important each dimension is for prediction: the shorter the length-scale, the more important the input (Rasmussen and Nickisch, 2016a).
- Due to the large size of the data sets, the Fully Independent Training Conditional (FITC) approximation was used for approximating the covariance func-

tion. The FITC approximation uses a low-rank diagonal matrix $\tilde{\mathbf{K}} = \mathbf{Q} + \text{diag}(\mathbf{K} - \mathbf{Q})$ where $\mathbf{Q} = \mathbf{K}_u^\top \mathbf{K}_{uu}^{-1} \mathbf{K}_u$, instead of the exact covariance matrix \mathbf{K} . \mathbf{K}_u and \mathbf{K}_{uu} contain covariances and cross-covariances for and between inducing points \mathbf{u}^i and data points \mathbf{x}^j (Rasmussen and Nickisch, 2010). In the implementation described, inducing points are chosen to be the full set of training set points (Rasmussen and Nickisch, 2016b).

- Because the factorial likelihood given in 4.16 is non-Gaussian, the posterior distribution, 4.17, over the latent variables is also non-Gaussian. As such, the latent distribution, 4.20, and the predictive distribution, 4.21, do not have analytical solutions (Nickisch and Rasmussen, 2008). Because the probit sigmoid likelihood is differentiable, the Laplace Approximation (LA) is used. LA approximates the posterior by a Gaussian centered at the mode with an appropriate curvature. Further details about LA may be found in (Rasmussen and Williams, 2006) and (Nickisch and Rasmussen, 2008).

4.5.3 Uganda Case Study

The GPC implementation described above is applied to the sparse 2012 electrification survey from Section 4.2.3, with \mathbf{x}_i describing geospatial building coordinates and y_i reflecting electrification status for a given household. In agreement with the logistic regression experiment described in 4.4.2, the survey was broken up into a training set with a random 70% of the original buildings and a test set with the remainder.

Fig. 4-11 shows results of the GP implementation and the trained model hyper-parameters; figure 4-12 shows the same model zoomed into Uganda’s capital city, Kampala. It is qualitatively evident that the model attributes higher electrification probability to areas where there are higher densities of electrified buildings in the training set. The sparse and clustered nature of the survey data set confer inference characteristics that reflect probabilistic electrification clusters scattered throughout the country.

Table 4.4: GP compared to other classifiers for the 2012 Uganda case study

Method	Data Set	Acc.	Prec.	Rec	F_1 Score
All Neg.	Train	87.65	N/A	N/A	N/A
LR (Unreg.)	Train	89.21	0.63	0.30	0.41
Gaus. Proc.	Train	93.31	0.75	0.69	0.72
All Neg.	Test	87.64	N/A	N/A	N/A
LR (Unreg.)	Test	88.99	0.63	0.26	0.37
Gaus. Proc.	Test	88.47	0.53	0.54	0.54

Table 4.4 compares common error metrics for a naive all-negative classifier, the GPC, and the logistic regression method from Section 4.4.2. All of the metrics presented reflect a symmetric loss function: the electrified status is assigned if the model’s probabilistic prediction is greater than or equal to 0.5, and the non-electrified status is assigned otherwise. Only the accuracy metric helps to inform model comparison

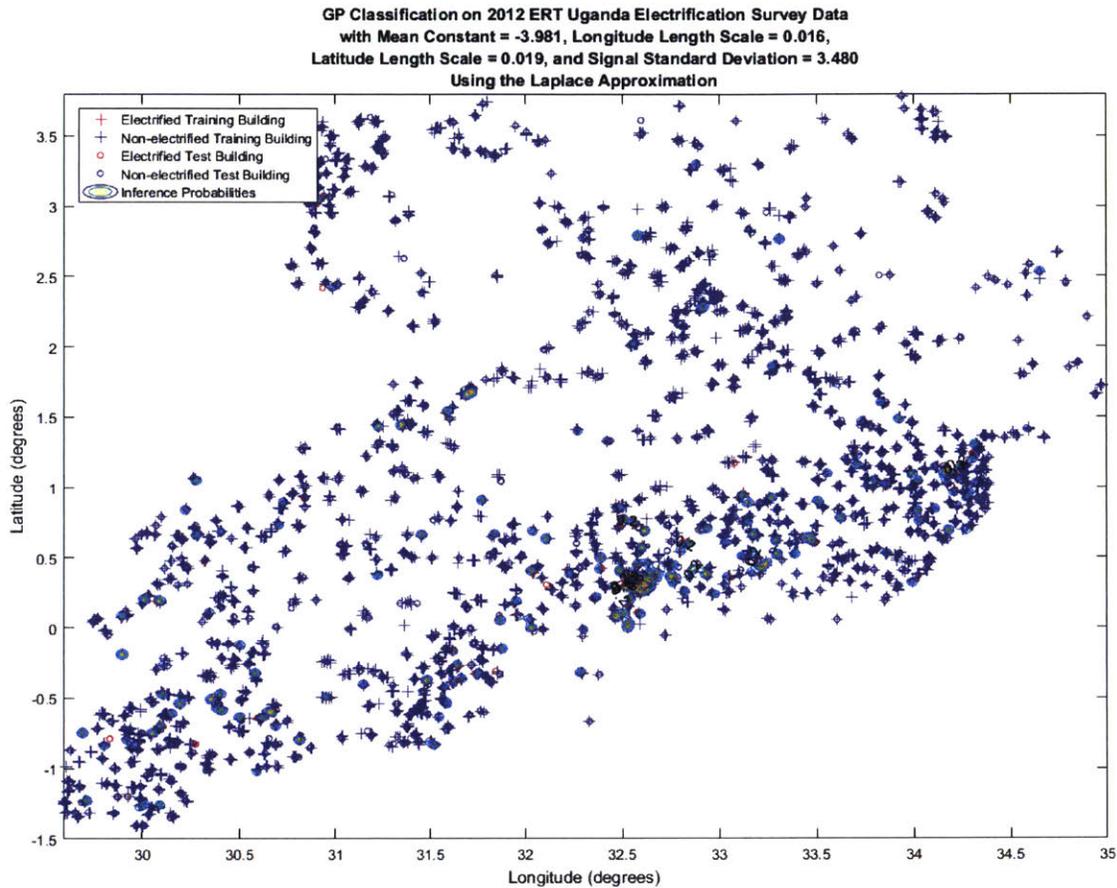


Figure 4-11: GP contour plots corresponding to electrification probabilities for all of Uganda using the 2012 ERT survey.

with the all-negative classifier, since the all-negative classifier does not assign positive values. Still, the all-negative classifier is the weakest performer with regards to accuracy, and the GPC approach produces significantly higher F_1 scores relative to the logistic regression model for both the training and test sets. The GPC's training and testing F_1 scores were 0.72 and 0.54 which the logistic regression scores were 0.41 and 0.37, respectively. While the logistic regression confers better test accuracy than the GPC approach, this is likely because it assigns more negative values, the majority class, as reflected with its higher precision and lower recall values.

Though the GPC F_1 scores are the highest of the models compared, the large discrepancy between them also suggests that significant overfitting is occurring. This makes sense given that the GPCs used are essentially just fitting a classifier around geospatial coordinates in the training set. In places where there is no training data, the GPC has no data to base its inference on, and simply interpolates the values of distant points. Nevertheless, the 0.54 test set F_1 score for GPC suggests that spatial correlation is an important feature for modeling electrification status. This conclusion is supported by intuition of network effects associated with connecting to

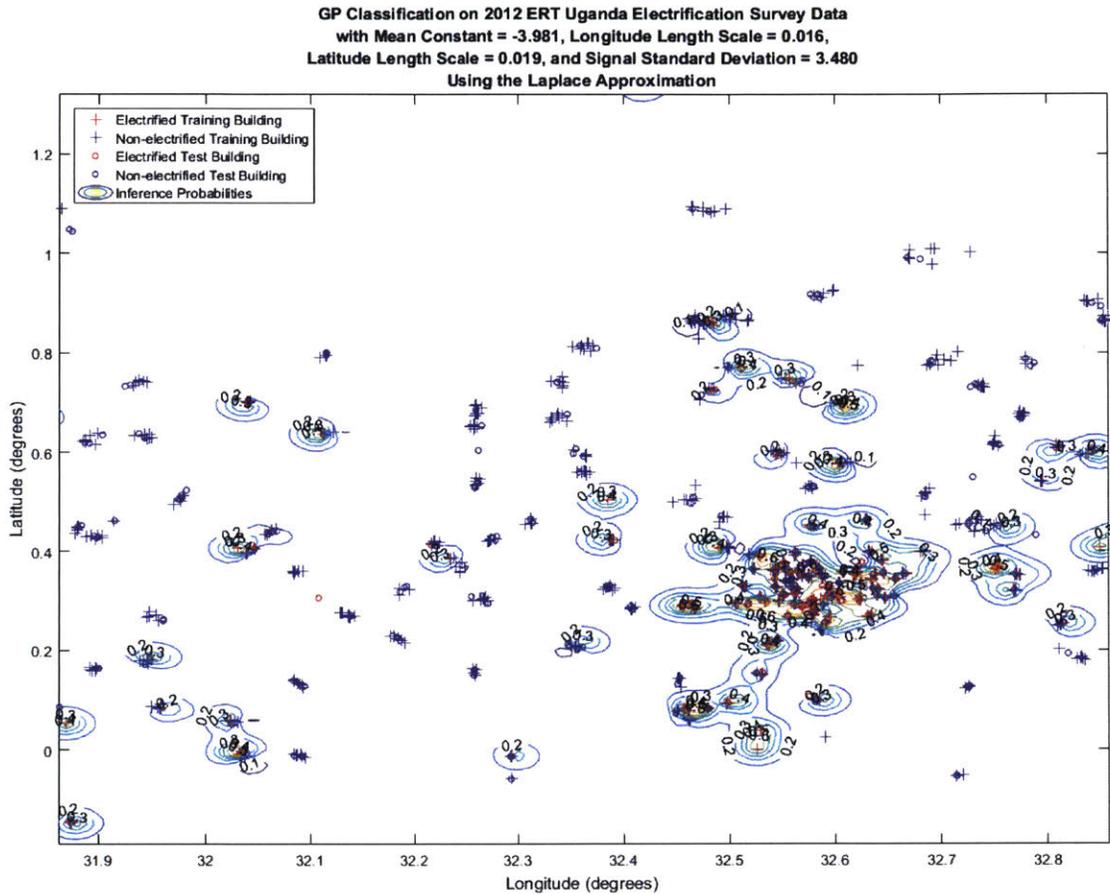


Figure 4-12: GP results corresponding to the 2012 ERT survey, zoomed into the capitol city, Kampala.

the low-voltage grid: it is much more economical to connect to the grid if a building is near to the existing grid than if it is far from it. Because grid-extension is the predominant mode of electrification in Uganda, electrified buildings are likely to be relatively close to other electrified buildings.

4.5.4 South Service Territory, Uganda Case Study

The same GP implementation was retrained and evaluated using the 2016 GIZ Uganda SST survey from Section 4.2.2. Results and trained hyperparameters are displayed and reported in figure 4-13 and table 4.5, respectively.

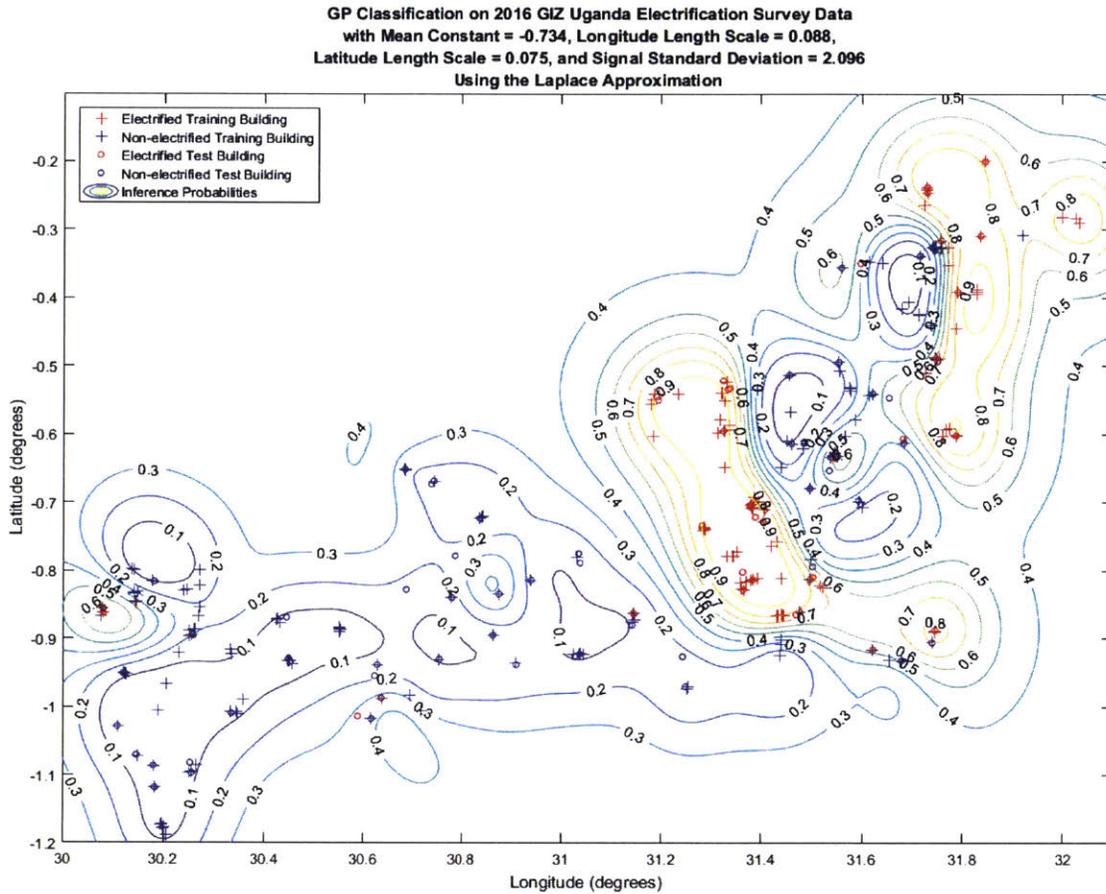


Figure 4-13: GP results corresponding to the 2016 REA survey for the Uganda SST

Table 4.5: GP compared to all-negative classifier for the 2016 Uganda SST case study

Method	Data Set	Acc.	Prec.	Rec	F_1 Score
All Neg.	Train	55.76	N/A	N/A	N/A
LR (Unreg.)	Train	57.58	1.00	0.04	0.08
Gaus. Proc.	Train	88.48	0.88	0.86	0.87
All Neg.	Test	55.63	N/A	N/A	N/A
LR (Unreg.)	Test	57.04	1.00	0.03	0.06
Gaus. Proc.	Test	85.92	0.84	0.84	0.84

Unlike the GP’s relative performance using the 2012 Uganda ERT survey, the GP classifier for the Uganda SST shows very large increases in classification accuracy relative to the all-negative classifier and to the logistic regression approach. This is likely be due to the fact that the share of electrified buildings in the 2016 survey is much greater than in the 2012 survey, resulting in a more even split between the majority and minority classes. The GP classifier’s improvement could also be due to the fact that the Uganda SST in 2016 had higher under-grid electrification rates than Uganda did at the country-level in 2012. The results here show the GP’s success

in exploiting high spatial correlation in electrification status, and how modeling this feature can be very informative.

4.5.5 Discussion

Despite the apparent success of the GP classifier approach over the logistic regression method for the cases presented in sections 4.5.3 and 4.5.4, the experimental design employed overlooks some of the GP classifier’s weaknesses. The nature of the survey data sets used for ground truth made it difficult to divide them into completely independent training and test sets. The survey points were collected of several buildings in loosely defined clusters spread across regions of interest. Partitioning these data sets using random sampling still results in training and test sets that are spatially correlated. This spatial correlation may affect the results from logistic regression, as neighboring buildings likely have similar nighttime light intensities, transformer distances, and electrification status. Nevertheless, it likely affects GPC even more, as having similar geospatial coordinates guarantees spatial proximity. Indeed, the GP classifier would be effectively incapable of making inferences in the absence of spatial correlation between training and test sets. In the simple formulation presented, it does not generalize well to areas far from its training set. This weakness, in addition to knowledge that spatial correlation is important for electrification status estimation, motivates the use of more flexible spatial models. We attempt to develop such a model in Section 4.6.

4.6 The Hierarchical Beta Model

The logistic regression and GP analyses presented in Sections 4.4 and 4.5 elucidate the need for spatial models with greater flexibility when considering electrification status estimation. Analyses using the GP classifier demonstrate that spatial correlation is an important indicator for electrification status; however, in the absence of nearby survey data, the GP classifier is severely limited. On the other hand, the logistic regression approach is able to capture more widely available features including nighttime light intensity and transformer distance, but it does not model spatial correlation in a direct way. Likely as a result of this, the logistic regression approach shows poorer classification performance relative to the GP classifier in both the Uganda and Uganda SST case studies. Being able to combine some of the benefits of each could help to improve our inferences across the board.

We further note that other features related to electrification are often available that neither the logistic regression method nor the GP classifier are well suited for. These include aggregate census metrics and incomplete building location information. For the Uganda SST, data regarding sub-county electrification rates are available from the 2014 UBOS census and can be geo-referenced to their appropriate administrative regions as described in Section 4.2.3. Such region-based data is not straightforward to incorporate into conventional methods. In addition, incomplete building location information cannot be easily accounted for either. In our data sets, we know that

the presence of a building coordinate means that a building is present; however, the absence of a building coordinate does not necessarily mean that a building is absent. For our Uganda SST case study, this is observed due to the fact that underlying satellite images for our building extractor were not available through the Google Maps API, as discussed in Sections 3.5 and 4.2.2.

In this section, we discuss ongoing progress towards the development of a hierarchical Bayesian model that is able to efficiently fuse measurements with different properties from multimodal sources at multiple scales. By dividing up areas of interest into grid cells and modeling electrification status as a set of latent variables over these cells, we are able to map region-based random variables to them and use measurements at both scales to infer electrification status. We simultaneously define mechanisms for enforcing spatial smoothness and take advantage of beta-binomial conjugacy to enable efficient inference. The model we present is designed to work under data sparse conditions which are common for the developing countries of interest. Finally, it can help provide characterizations of uncertainty that are useful for decision-making. We refer to this novel model architecture as the "hierarchical beta model."

The hierarchical beta model was developed in close collaboration with Christopher Dean and Dr. John Fisher at the MIT Computer Science and Artificial Intelligence Laboratory. Significant credit should be given to them and their ongoing research regarding spatial modeling and information planning. The general framework presented here may be additionally applicable in a number of different domains related to geospatial data fusion and inference.

4.6.1 Background

Very brief technical and methodological overviews are presented to provide context for the machine learning models proposed. Common probability distributions, Bayesian inference, probabilistic graphical models, sampling methods, hierarchical models, and empirical Bayes are discussed along with supporting topics. The purpose of these technical overviews is primarily contextual and as such, they should not be seen as comprehensive treatments. Excellent references for further reading on probabilistic graphical models and machine learning in general include (Koller and Friedman, 2009; Bishop, 2006; Murphy, 2012).

Bayes' Rule

Bayes' rule or Bayes' theorem, given in 4.23, enables Bayesian inference. The Gaussian process and Bayesian network models presented in this chapter exemplify its use.

$$p_{x|y}(x|y) = \frac{p_x(x)p_{y|x}(y|x)}{p_y(y)} = \frac{p_x(x)p_{y|x}(y|x)}{\sum_{x'} p_x(x')p_{y|x}(y|x')} \quad (4.23)$$

where $p_x(x)$ and $p_y(y)$ represent the prior distributions, while $p_{y|x}(y|x)$ denotes the likelihood and $p_{x|y}(x|y)$ corresponds to the posterior distribution of x given y . In

essence, the prior represents “prior knowledge” of a distribution of interest, and the posterior represents updated beliefs about the distribution given data or observations.

The Binomial Distribution

While the Bernoulli distribution as presented in 4.4.1 describes only single binary random variables, the Binomial distribution is a pmf that concerns n binary random variables. Consider a data set $\mathcal{D} = \{x_1, \dots, x_n\}$ composed of observed values of $x_i \in \{0, 1\}$ with probability of ‘success’ θ with $\theta \in [0, 1]$. The Binomial distribution is a distribution over the number of successes $x \in \{0, 1, \dots, n\}$ with pmf:

$$p_x(k) = \Pr\{x = k\} = \text{Binomial}(k; n, \theta) \triangleq \binom{n}{k} \theta^k (1 - \theta)^{n-k} \quad (4.24)$$

In this distribution k represents the number of instances of $x_i = 1$ observed given n trials. This distribution has mean and variances given by 4.25 and 4.26.

$$\mathbb{E}[x] = n\theta \quad (4.25)$$

$$\text{var}(x) = n\theta(1 - \theta) \quad (4.26)$$

The Beta Distribution

The beta distribution is a distribution over the interval $[0, 1]$ with probability density function (pdf) as given by 4.27. The beta distribution is a conjugate prior for the binomial distribution and is often used to represent uncertainty about the success probability θ .

$$p_x(x) = \text{Beta}(x; \alpha, \beta) \triangleq \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1} (1 - x)^{\beta-1} \quad (4.27)$$

The beta distribution’s normalizing coefficient, given by the ratio of gamma functions¹, ensures that distribution is normalized (its integral over $x \in [0, 1]$ is equal to one). The hyperparameters α and β control the distribution of x . The mean, variance, and mode of the beta distribution are defined by 4.28, 4.29, and 4.30, respectively.

$$\mathbb{E}[x] = \frac{\alpha}{\alpha + \beta} \quad (4.28)$$

$$\text{var}(x) = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)} \quad (4.29)$$

¹ The gamma function is defined by

$$\Gamma(x) \triangleq \int_0^\infty u^{x-1} e^{-u} du.$$

$$\text{mode}(x) = \frac{\alpha - 1}{\alpha + \beta - 2} \quad (4.30)$$

Plots of the beta distribution for various hyperparameter values are shown in Fig. 4-14a. Note that when $\alpha = \beta = 1$, we obtain the uniform distribution.

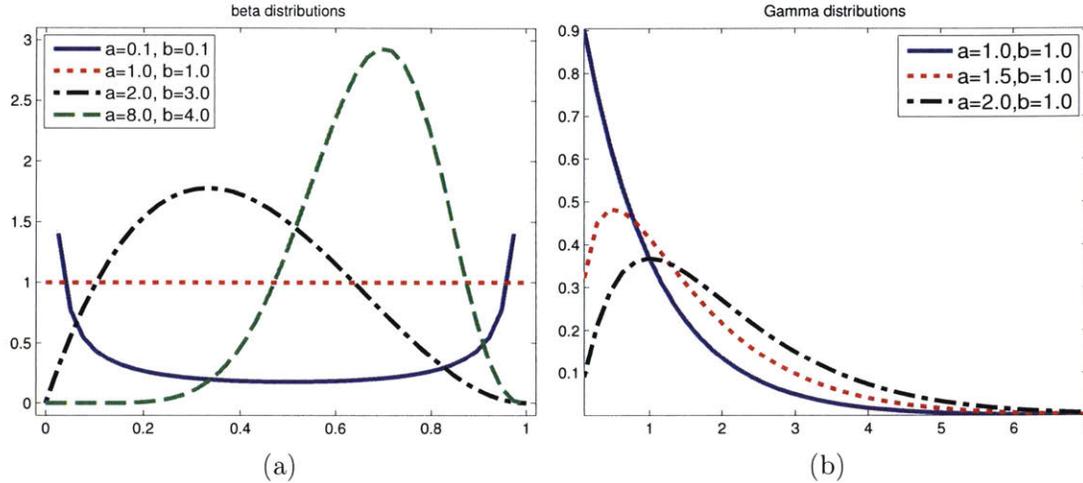


Figure 4-14: (a) Beta distribution examples. (b) Gamma distribution examples. Figure sources: (Murphy, 2012)

The beta-binomial model serves as a component of the hierarchical Bayes model to be presented. If a beta distribution is used as a prior on θ , it is conjugate with binomial observations and yields a beta posterior distribution. Conjugacy significantly simplifies computation. Specifically, given a beta prior with the form $p(\theta; \alpha, \beta) \propto \theta^{\alpha-1}(1-\theta)^{\beta-1}$ and a binomial likelihood, $p(\mathcal{D}|\theta) = \binom{l+m}{m} \theta^m (1-\theta)^l$, with m ‘heads’ and l ‘tails,’ we use Bayes rule to compute a beta posterior: $p(\theta|\mathcal{D}) \propto \theta^{m+\alpha-1}(1-\theta)^{l+\beta-1}$. The posterior can be written another way as $\theta|\mathcal{D} \sim \text{Beta}(\alpha + m, \beta + l)$. As such, α and β are interpreted as pseudocounts, because they affect the posterior in the same way as empirical counts from the binomial likelihood.

The Gamma Distribution

The Gamma distribution is a distribution over the positive real numbers. Note that the gamma distribution and the gamma function are *not* the same! A gamma distributed random variable x has pdf given by 4.31 where $x > 0$ with parameters $a > 0$, referred to as the shape, and $b > 0$, referred to as the rate.

$$p_x(x) = \text{Gamma}(x; a, b) \triangleq \frac{1}{\Gamma(a)} b^a x^{a-1} \exp(-bx) \quad (4.31)$$

The mean and variance of the gamma distribution are given below by 4.32 and 4.33. If $a \leq 1$ the mode is 0. The mode is positive and equal to $(a - 1)/b$ otherwise. Plots of the distribution for different values of a and b are shown by Fig. 4-14b.

$$\mathbb{E}[\mathbf{x}] = \frac{a}{b} \quad (4.32)$$

$$\text{var}(\mathbf{x}) = \frac{a}{b^2} \quad (4.33)$$

Probabilistic Graphical Models

Probabilistic graphical models help to solve the problem of intractability while representing a joint distribution $p_{\mathbf{X}}$ over a large set of random variables $\mathbf{X} = \{x_1, \dots, x_N\}$. As (Koller and Friedman, 2009) illustrates, representing a joint distribution can quickly become unmanageable for even the simplest case of binary-valued random variables. In this case, \mathbf{X} has 2^N possible assignments, requiring $p_{\mathbf{X}}$ to be represented with $2^N - 1$ parameters. Very large numbers of parameters are infeasible to process computationally, cognitively, and statistically. Probabilistic graphical models constitute a framework to cut through this complexity by exploiting the representation of independence properties within the joint distribution. Nodes encode subsets of random variables and edges depict statistical relationships between the nodes. The graphical models we discuss allow representation, inference, and learning to support decision-making under uncertainty.

Directed Graphical Models

Directed graphical models, also known as Bayesian networks, Bayes nets, or belief networks, provide a representation of the independence properties associated with a distribution. A Bayesian network is represented by a directed acyclic graph (DAG), $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where nodes $\mathcal{V} = \{1, 2, \dots, N\}$ represent random variables and directed edges are given as $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$. The notation $(i, j) \in \mathcal{E}$ denotes an edge from node i to node j .

Bayesian networks define families of distributions that factor by functions of nodes and their parents. Specifically, the functions $f_i(x_i, x_{\pi_i})$ represent probability distributions for x_i conditioned on its parent variables x_{π_i} , as shown in 4.34.

$$p_{x_1, \dots, x_N}(x_1, \dots, x_N) = \prod_{i=1}^N f_i(x_i, x_{\pi_i}) = \prod_{i=1}^N p_{x_i | x_{\pi_i}}(x_i | x_{\pi_i}) \quad (4.34)$$

The functions $f_i(x_i, x_{\pi_i})$ must be non-negative-valued and for any given set of parent values, must sum to one.

$$\sum_{x_i \in \mathcal{X}} f_i(x_i, x_{\pi_i}) = 1 \quad (4.35)$$

Directed graphical model representations provide a framework to support inference and learning.

Inference

Within the subject of inference over probabilistic graphical models, we discuss the core task of calculating posterior beliefs. When calculating posteriors, in general, we model $p_{\mathbf{x},\mathbf{y},\mathbf{z}}(\mathbf{x},\mathbf{y},\mathbf{z})$, where \mathbf{x} represents latent variables, \mathbf{y} represents observed variables, and \mathbf{z} represents latent variables not of direct interest. Considering a prior belief of $p_{\mathbf{x}}(\mathbf{x})$, we seek to calculate the posterior represented by 4.36. Note that for continuous distributions, summations should be replaced with integrations.

$$p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y}) = \frac{p_{\mathbf{x},\mathbf{y}}(\mathbf{x},\mathbf{y})}{p_{\mathbf{y}}(\mathbf{y})} = \frac{\sum_{\mathbf{z}} p_{\mathbf{x},\mathbf{y},\mathbf{z}}(\mathbf{x},\mathbf{y},\mathbf{z})}{\sum_{\mathbf{x},\mathbf{z}} p_{\mathbf{x},\mathbf{y},\mathbf{z}}(\mathbf{x},\mathbf{y},\mathbf{z})} \propto \sum_{\mathbf{z}} p_{\mathbf{x},\mathbf{y},\mathbf{z}}(\mathbf{x},\mathbf{y},\mathbf{z}) \quad (4.36)$$

For many models, it is infeasible to evaluate the posterior distribution or even compute expectations using exact inference for a variety of reasons. For continuous variables, the required integrations may not have closed-form analytical solutions and high dimensionality may preclude numerical integration. For discrete variables, requisite summations may be too computationally expensive to do in practice.

Instead of exact inference, methods for approximate inference are often pursued. Approximate inference methods fall into two classes depending on whether they make stochastic or deterministic approximations. Deterministic methods such as variational inference are based on producing analytical approximations for the posterior distribution. While they are often used to put bounds on the posterior distribution, they generally do not produce exact results and will not be expounded on here. On the other hand, stochastic techniques such as Markov chain Monte Carlo (MCMC) rely on sampling from the posterior distribution. They converge to exact results given infinite computational resources; however, with finite computation time, they provide only approximate solutions. Since MCMC methods are featured in the hierarchical beta model presented, we describe them in more detail in the following sections.

Sampling and Markov Chain Monte Carlo Methods

Sampling methods, also known as Monte Carlo methods, are used to characterize distributions based off of samples. Generally, the idea is that if we have enough samples, we can learn anything we want about a distribution. This intuition is portrayed by the strong law of large numbers, which states that

$$\frac{1}{K} \sum_{k=1}^K f(\mathbf{x}^k) \xrightarrow{\text{a.s.}} \mathbb{E}[f(\mathbf{x})] \quad \text{when } k \rightarrow \infty \quad (4.37)$$

where $\xrightarrow{\text{a.s.}}$ represents convergence with probability one. Choosing different functions f enables the quantification of any aspect of p . Setting $f(\mathbf{x}) = \mathbf{x}$ gives the mean, while $f(\mathbf{x}) = (\mathbf{x} - \mathbb{E}[\mathbf{x}])^2$ conveys the variance, $f(\mathbf{x}) = -\log p(\mathbf{x})$ the entropy, and $f(\mathbf{x}) = \mathbb{1}_{\mathbf{x} \in \mathcal{S}}$ the probability $p(\mathbf{x} \in \mathcal{S})$. Marginals can also be computed from samples simply by dropping values of variables that are not of direct interest. For instance, if $\mathbf{x}^1, \dots, \mathbf{x}^K$ are samples from the joint distribution $p(\mathbf{x})$, then $\mathbf{x}_i^1, \dots, \mathbf{x}_i^K$ are samples from the marginal distribution $p(x_i)$.

Monte Carlo methods allow us to sample from complex distributions for which we can only evaluate up to a normalizing constant, Z . Suppose we want to sample from $p(x)$,

$$p(x) = \frac{\tilde{p}(x)}{Z}, \quad x \in \mathcal{X}, \quad (4.38)$$

where we can evaluate $\tilde{p}(x)$ but where \mathcal{X} is sufficiently large to make Z infeasible to compute. While Monte Carlo methods such as rejection sampling enable us to produce samples of $p(x)$, they tend to be inefficient when \mathcal{X} is large. Ironically, this is exactly the situation when we usually resort to sampling in the first place.

MCMC is a related method for sampling from distributions which scales well with large \mathcal{X} and high dimensionality. As with rejection sampling, MCMC samples from a proposal distribution; however, it circumvents some of the associated constraints on the sampling procedures which render rejection sampling inefficient. The key behind MCMC is to construct a Markov chain whose stationary distribution is $p(\mathbf{x})$. We essentially perform a random walk on the state space such that the fraction of time spent at any given state \mathbf{x} is equal to $p(\mathbf{x})$.

Metropolis Hastings Algorithm

The Metropolis Hastings (MH) algorithm is a general algorithm for constructing a Markov chain of samples from p as presented in our description of MCMC. At each step in MH, we move to state \mathbf{x}' given information about the current state \mathbf{x}^s , and thus generate a new sample. MH does this relying on a proposal distribution $q(\mathbf{x}'|\mathbf{x})$ from which we can sample directly, using it to define an acceptance probability, and sampling from a uniform distribution to determine whether or not to accept the proposal as a sample. The Metropolis Hastings algorithm, outlined in Algorithm 3, begins by defining an arbitrary starting point x^0 where $p(x^0) > 0$, and iterating through S samples.

The number of iterations S must be large enough to ensure adequate mixing time so that the Markov chain has "burned in" and samples from the stationary distribution. For this reason, it is often necessary to discard some of the initial samples. The collection of a sufficient number of samples is also required to converge on the distribution's metrics of interest by the strong law of large numbers. Subsampling can also be pursued to decrease dependency between samples.

A proposal distribution that is often used is the symmetric Gaussian centered on the current state $q(\mathbf{x}'|\mathbf{x}) = \mathcal{N}(\mathbf{x}'|\mathbf{x}, \Sigma)$; when MH is used with this proposal distribution, it is referred to as the random walk Metropolis Hastings algorithm. Σ is generally tuned to keep the rejection rate low while allowing adequate coverage of the state space. If the proposal distribution for MH is symmetric, as in the random walk Metropolis Hastings algorithm, $q(\mathbf{x}'|\mathbf{x}) = q(\mathbf{x}|\mathbf{x}')$ and the acceptance probability becomes

$$\alpha = \frac{\tilde{p}(x')}{\tilde{p}(x)}.$$

Algorithm 3 Metropolis Hastings algorithm

- 1: **procedure** METROPOLISHASTINGS(x^0, S)
- 2: **for** $s = 0, 1, 2, \dots, S - 1$ **do**
- 3: Define $x = x^s$;
- 4: Sample $x' \sim q(x'|x)$;
- 5: Compute acceptance probability

$$\alpha = \frac{\tilde{p}(x')q(x|x')}{\tilde{p}(x)q(x'|x)}$$

- 6: Compute $r = \min(1, \alpha)$
- 7: Sample $u \sim \text{Uniform}(0, 1)$
- 8: Set new sample

$$x^{s+1} = \begin{cases} x' & \text{if } u < r \\ x^s & \text{if } u \geq r \end{cases} .$$

return $\{x^s\}_{s=1}^S$

Gibbs Sampling

Gibbs sampling, a special case of MH, is presented in Algorithm 4 and directly exploits the factorization structure of a target distribution. The basic idea is to iterate through variables and sample each of them conditioned on all of the other variables in the distribution. For sparse graphical models, conditional independencies generally simplify the conditional distributions.

Algorithm 4 Gibbs sampler algorithm

- 1: **procedure** GIBBSAMPLER(x^0, S)
 - 2: **for** $s = 0, 1, 2, \dots, S - 1$ **do**
 - 3: Sample $x_1^{s+1} \sim p(x_1|x_2^s, x_3^s, \dots, x_M^s)$
 - 4: Sample $x_2^{s+1} \sim p(x_2|x_1^{s+1}, x_3^s, \dots, x_M^s)$
 - 5: \vdots
 - 6: Sample $x_j^{s+1} \sim p(x_j|x_1^{s+1}, \dots, x_{j-1}^{s+1}, x_{j+1}^s, \dots, x_M^s)$
 - 7: \vdots
 - 8: Sample $x_M^{s+1} \sim p(x_M|x_1^{s+1}, x_2^{s+1}, \dots, x_{M-1}^{s+1})$
- return** $\{x^s\}_{s=1}^S$
-

As in MH, initial samples from the Markov chain must be discarded to make sure that it has reached its stationary distribution. Furthermore, successive samples from the Markov chain will be highly correlated, making it beneficial to subsample the sequence to achieve more independent samples. Justifications for why Gibbs sampling is a special case of MH are given in (Murphy, 2012) and (Dean, 2015).

Hierarchical Bayes

Hierarchical Bayesian models, also called multi-level models, involve treating unknown hyperparameters as hidden variables; we refer to hyperparameters as parameters of the prior distribution. This is an alternative to setting uninformative priors and can enable random variables lower in a hierarchy to influence one another. For example, we can turn a one-level model $\boldsymbol{\theta} \rightarrow \mathcal{D}$, where $\boldsymbol{\theta}$ has hyperparameters $\boldsymbol{\eta}$, into a two-level model treating $\boldsymbol{\eta}$ as a hidden variable:

$$\boldsymbol{\eta} \rightarrow \boldsymbol{\theta} \rightarrow \mathcal{D}.$$

With multiple θ_i parameterized by the same $\boldsymbol{\eta}$, a value θ_i with fewer observations can borrow statistical strength from $\boldsymbol{\theta}_{-i}$ through their shared hyperparameters, $\boldsymbol{\eta}$. This can be seen as a compromise between estimating values of each θ_i separately and between assuming that all θ_i are equivalent. The former is problematic when there may be too little evidence to support a given θ_i , and the latter may result in model oversimplification.

Empirical Bayes

Empirical Bayes entails an approximation to using hierarchical Bayesian models. In the two-level model presented in Section 4.6.1, we see that we need to compute:

$$p(\boldsymbol{\eta}, \boldsymbol{\theta} | \mathcal{D}) \propto p(\mathcal{D} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \boldsymbol{\eta}) p(\boldsymbol{\eta}).$$

In some cases, it is possible to analytically marginalize out $\boldsymbol{\theta}$ and $p(\boldsymbol{\eta} | \mathcal{D})$ can be computed; however, this is not always feasible. Instead, we can approximate $\boldsymbol{\eta}$ with point estimate $p(\boldsymbol{\eta} | \mathcal{D}) \approx \delta_{\hat{\boldsymbol{\eta}}}(\boldsymbol{\eta})$, where $\hat{\boldsymbol{\eta}} = \operatorname{argmax}_{\boldsymbol{\eta}} p(\boldsymbol{\eta} | \mathcal{D})$. Using a uniform prior on $\boldsymbol{\eta}$ allows us to simplify this integral into the following:

$$\hat{\boldsymbol{\eta}} = \operatorname{argmax}_{\boldsymbol{\eta}} \left[\int p(\mathcal{D} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \boldsymbol{\eta}) d\boldsymbol{\theta} \right] = \operatorname{argmax}_{\boldsymbol{\eta}} p(\mathcal{D} | \boldsymbol{\eta}) \quad (4.39)$$

While empirical Bayes conflicts with the Bayesian principle that priors should be chosen independently of the data, it provides an approximation to inference on a hierarchical Bayesian model and has shown to be effective in a variety of applications (Casella, 1985; Murphy, 2012).

4.6.2 Methods

A general depiction of the hierarchical beta model is given in Fig. 4-15. Geospatial areas of interest are divided into an arbitrarily sized two-dimensional regular grid with latent variables representing beta distributions over an attribute of interest for each cell. Region-level latent variables also take on beta distributions and represent aggregated attributes for groups of cells. Regions map to one or more cells and cells may belong to zero or more regions. Sparse observations are provided at both the region and cell levels and take on any data format that can be converted into ‘counts’

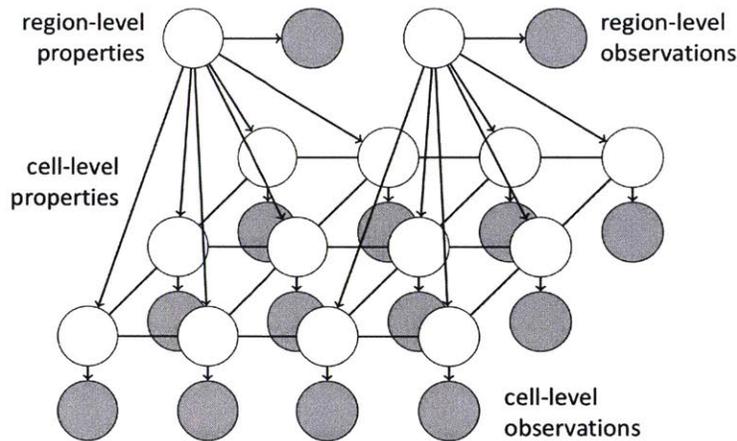


Figure 4-15: Visualization of a hierarchical beta model. White nodes are latent beta variables and shaded nodes are observed values modeled as binomials.

for beta-binomial updates. Additional latent variables are incorporated to promote spatial and multiscale smoothness and sampling algorithms are employed to evaluate posterior distributions.

When applying the hierarchical beta model to electrification status estimation, we define electrification probability to be the attribute of interest corresponding to each region and cell latent variable. Observations we present at the cell-level include annual composite nighttime light intensity, sparse survey data, transformer distance, transformer count, and building density. Observations at the region level correspond to aggregate census electrification metrics.

Fig. 4-16 provides a representation of the hierarchical beta model using plate notation and equations 4.40 through 4.43 define latent variable priors, observation likelihoods, and the joint probability distribution. The model has M region and L edge plates mapping to N cell plates. Edges between plates define region assignments and cell connectivity along a regular grid. In the following sections, we discuss how various components fit together and describe key features including multiscale smoothness, spatial smoothness, efficient inference, function approximation, and posterior analysis.

Enforcing Multiscale Smoothness

Multiscale smoothness refers to consistency between region-level features and associated cell-level features. This allows beliefs about regions to affect cell beliefs and vice versa. In addition, cell beliefs can influence one another through connections to common region variables. Likewise, regions beliefs can influence one another to a limited extent through connections to common cell variables. Since we use census-mapped administrative boundaries to define regions, cell variables only have connections to multiple region variables if they correspond to areas within multiple regions.

Specifically, each region plate has a beta random variable μ_j denoting region electrification status. Hyperparameters α_μ and β_μ are common to all regions variables

$$\theta_i | \boldsymbol{\mu}, \boldsymbol{\gamma}, \boldsymbol{\varepsilon} \sim \text{Beta}(\underbrace{\gamma_\mu \mathbf{w}_i^\top \boldsymbol{\mu} + \gamma_\varepsilon \text{mean}_{k \in \mathcal{N}(i)} \varepsilon_k}_{\alpha_{\theta_i}}, \underbrace{\gamma_\mu (1 - \mathbf{w}_i^\top \boldsymbol{\mu}) + \gamma_\varepsilon (1 - \text{mean}_{k \in \mathcal{N}(i)} \varepsilon_k)}_{\beta_{\theta_i}}) \quad (4.43)$$

Cell prior

The formulation presented for multiscale smoothness can make sense for the application of electrification status estimation because populations within a given small-scale administrative region tend to have similar socioeconomic status on average and similar infrastructure characteristics. Aggregate region-level metrics on electrification status may be representative of a region’s constituent cells in the absence of indicators suggesting otherwise. Conversely, strong evidence at the cell-level should naturally affect beliefs at the region-level, which by definition are aggregate characterizations of the lower levels.

There is opportunity for future work on region-cell connectivity to enhance the representation provided here. Because region metrics are generally derived from censuses, it may be more appropriate to treat them as capacity constraints on cell-level electrification. Combined with data regarding the number of buildings in a region, region-level electrification probabilities can be used to compute expected numbers of electrified and non-electrified buildings. Modifications that allow a given region variable to polarize cell electrification status while matching these numbers in aggregate may confer improved performance. Further work must be done to compare multiscale capacity constraints with the current implementation’s approach of biasing all cells towards their mean electrification status.

Enforcing Spatial Smoothness

Spatial smoothness is enabled through the definition of edge variables, ε_k , connecting neighboring cell variables. Edge variables are beta distributed and have shared hyperparameters α_ε and β_ε as given by 4.44. Similar to the design of multiscale smoothness, spatial smoothness also has components that affect the hyperparameters of their associated θ_i as reflected in 4.43. These edge components bias θ_i towards the mean of their neighboring edge variables and the strength of this effect is controlled by the gamma distributed edge stickiness variable γ_ε as shown in 4.45. As with the region stickiness variable, γ_ε has its own hyperparameters $\alpha_{\gamma_\varepsilon}$ and $\beta_{\gamma_\varepsilon}$.

$$\varepsilon_k \sim \text{Beta}(\alpha_\varepsilon, \beta_\varepsilon) \quad (4.44)$$

Edge prior

$$\gamma_\varepsilon \sim \text{Gamma}(\alpha_{\gamma_\varepsilon}, \beta_{\gamma_\varepsilon}) \quad (4.45)$$

Edge stickiness prior

Fig. 4-17 portrays different one-dimensional representations of the hierarchical beta model. Fig. 4-17a depicts each latent variable and observation type discussed, including edge variables as parents of neighboring cell variables. On the other hand, 4-17b shows a representation of the same graph with edge variables marginalized out.

This marginalization demonstrates that the net effect of edge variables is inducing spatial similarity among cell variables corresponding to neighboring cells; nodes that share a common parent edge variable can be represented as if they are connected by an equivalent undirected edge in the marginalized graphical representation.

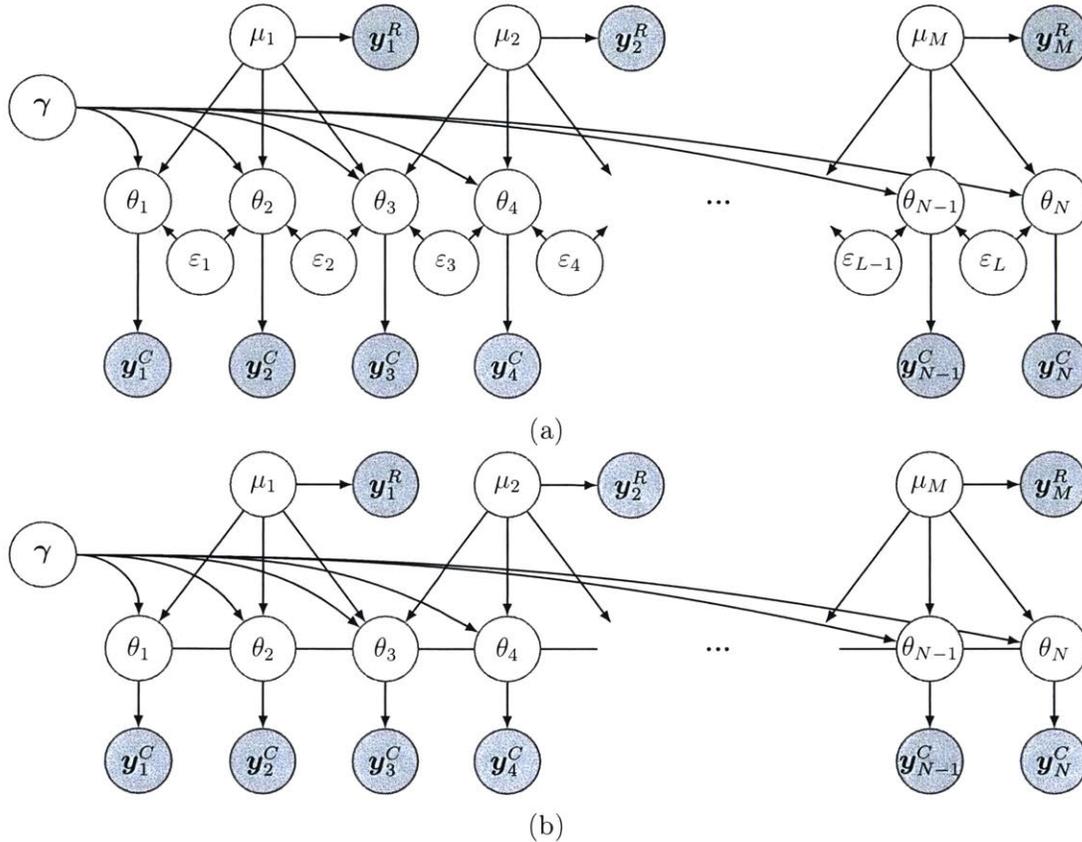


Figure 4-17: (a) A one-dimensional view of the hierarchical beta model with an extended graphical representation. Hyperparameters are not shown. (b) Edge variables are marginalized out, showing equivalent undirected edges between adjacent cell nodes.

A series of smoothness tests using synthetic data are presented in Fig. 4-18. In each frame, the top left and bottom right cells are pegged to low and high electrification status values, respectively. Node marginals from the posterior distribution are shown for all of the cells in each frame. Smoothness is evaluated as γ_ε increases from left to right starting with the top row. We see that while some γ_ε values induce visually appealing representations of smoothness, values that are too high reflect non-smoothness. This may be due to effects that γ_ε has on mixing times when using MCMC methods for inference; high values of γ_ε are seen to dramatically slow convergence. Further research is needed to improve computationally efficient representations of spatial smoothness over a range of γ_ε values.

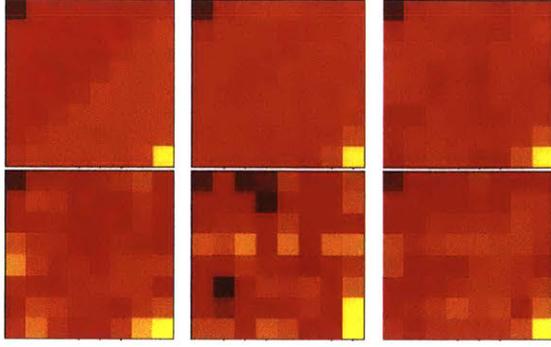


Figure 4-18: Node marginals for a series of synthetic test cases after inference with varying γ_ε . Values of γ_ε increase by orders of magnitude from the top left to the top right plot, and from the bottom left to the bottom right.

Function Approximation with Empirical Bayes

Just as region-level observations have binomial likelihoods given μ_j as described in 4.41, cell-level observations also have binomial likelihood given θ_i as in 4.46. All together, we define the joint distribution of the model by 4.47.

$$\mathbf{y}_i^C | \theta_i \sim \text{Binomial}(\mathbf{N}_i^C, \theta_i) \quad (4.46)$$

$$p(\boldsymbol{\mu}, \boldsymbol{\varepsilon}, \boldsymbol{\gamma}, \boldsymbol{\theta}, \mathbf{y}^R, \mathbf{y}^C) = p(\gamma_\mu) p(\gamma_\varepsilon) \prod_{k=1}^L p(\varepsilon_k) \prod_{j=1}^M \left[p(\mu_j) \prod_{k=1}^{D^R} p(y_{jk}^R | \mu_j) \right] \prod_{i=1}^N \left[p(\theta_i | \boldsymbol{\mu}, \boldsymbol{\varepsilon}, \boldsymbol{\gamma}, \mathbf{w}) \prod_{k=1}^{D^C} p(y_{ik}^C | \theta_i) \right] \quad (4.47)$$

For both region and cell-level observations, we can exploit the empirical Bayes formulation described in Section 4.6.1 to learn functions mapping input data types to counts for beta-binomial updates as long as we have one direct observation type at both the region and cell scales. As a first approximation, we also make the assumption that cell nodes are independent from region and edge nodes. We will later describe how we can relax this assumption and review the implications of this approach.

We only present the empirical Bayes procedure for cell-level observations as the case studies we present for the Uganda SST and Kayonza lack multimodal inputs at the region-level. For region-level inputs, we assume census metrics other than electrification status are conditionally independent of region latent variables given complete electrification status information. Though the following notation only refers to cell-level empirical Bayes, the methods and discussion are equivalent if we were to extend to the region level.

$$\theta_i \sim \text{Beta}(\alpha_i, \beta_i) \quad (4.48)$$

$$y_i^s \sim \text{Binomial}(\theta_i; n_i^s) \quad (4.49)$$

$$k(x_i^d; \phi^d) \sim \text{Binomial}(\theta_i; n(x_i^d; \phi^d)) \quad \forall d \in 1, 2, \dots, D \quad (4.50)$$

In this restricted formulation, cell nodes θ_i still have beta distributions but are independent of one another. Instead of being conditioned on μ , γ , and ϵ as in 4.43, we now define hyperparameters α_i and β_i as in 4.48. One observation type available for some of the cells must represent direct binomial observations from θ_i given by y_i^s as defined in 4.49. For the electrification status application at the cell-level, y_i^s represents the number of electrified buildings surveyed within a cell while n_i^s represents the total number of buildings surveyed in the cell. Functions relating D other observations types to count data can modeled as given by 4.50 and learned using empirical Bayes. These functions represent type-specific transformations. Here x_i^d represents a numerical value in the native feature's scaling while $k(x_i^d; \phi^d)$ represents the corresponding number of 'success' counts and $n(x_i^d; \phi^d)$ represents the number of total counts. We decompose these functions according to 4.51 and 4.52, where $a(x_i^d; \phi^d)$ equivalently implies the total number of 'success' counts while $b(x_i^d; \phi^d)$ explicitly measures 'failure' counts.

$$k(x_i^d; \phi^d) = a(x_i^d; \phi^d) \quad (4.51)$$

$$n(x_i^d; \phi^d) = a(x_i^d; \phi^d) + b(x_i^d; \phi^d) \quad (4.52)$$

Restricting the type-specific transformations to linear functions for simplicity, we describe equations 4.53 and 4.54 and seek to learn parameter values in ϕ^d for all D data types. Non-linear transformations may also be learned by modifying 4.53 and 4.54 and their associated optimization constraints, which will be described.

$$a(x_i^d; \phi^d) = \phi_1^d x_i^d + \phi_2^d \quad (4.53)$$

$$b(x_i^d; \phi^d) = \phi_3^d x_i^d + \phi_4^d \quad (4.54)$$

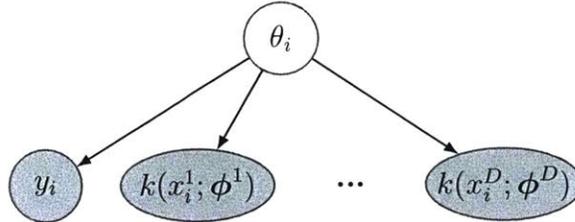


Figure 4-19: Empirical Bayes in the hierarchical beta model.

Empirical Bayes is applied to approximate ϕ as represented by Fig. 4-19. We

formulate the problem analogously to the two-level model presented in Section 4.6.1, where data \mathcal{D} is composed of the set of direct observations for individual cells y_i ; however, we set the parameter distribution to be the posterior distribution after observing all of the D transformed observation types, $p(\theta_i|\phi, \{x_i^d\})$. We represent the approximation with 4.55 where direct observation probabilities are given by 4.56 and our posterior distribution after observing the transformed observation types is distributed as 4.57.

$$\phi^* = \operatorname{argmax}_{\phi^1, \phi^2, \dots, \phi^D} \prod_{i=1}^N \int p(y_i^s|\theta_i) p(\theta_i|\phi, \{x_i^d\}) d\theta_i \quad (4.55)$$

$$p(y_i^s|\theta_i) = \text{Binomial}(y_i^s|\theta_i; n_i^s) \quad (4.56)$$

$$p(\theta_i|\phi, \{x_i^d\}) = \text{Beta}(\theta_i|\alpha_i + \sum_{d=1}^D a(x_i^d; \phi^d), \beta_i + \sum_{d+1}^D b(x_i^d; \phi^d)) \quad (4.57)$$

We note that the integral in 4.55 is distributed according to the beta-binomial distribution, an analytically tractable compound distribution. We equate this integral to $p(y_i^s|\phi, \{x_i^d\})$ in 4.58, redefine its hyperparameters $\alpha'_i(\phi)$ and $\beta'_i(\phi)$ in 4.59, and provide an explicit closed-form representation in 4.60.

$$p(y_i^s|\phi, \{x_i^d\}) = \int p(y_i^s|\theta_i) p(\theta_i|\phi, \{x_i^d\}) d\theta_i \quad (4.58)$$

$$= \text{BetaBinomial}\left(\underbrace{y_i^s|\alpha_i + \sum_{d=1}^D a(x_i^d; \phi^d)}_{\alpha'_i(\phi)}, \underbrace{\beta_i + \sum_{d+1}^D b(x_i^d; \phi^d)}_{\beta'_i(\phi)}; n_i^s\right) \quad (4.59)$$

$$= \frac{\Gamma(n_i^s + 1)}{\Gamma(y_i^s + 1)\Gamma(n_i^s - y_i^s + 1)} \frac{\Gamma(y_i^s + \alpha'_i(\phi))\Gamma(n_i^s - y_i^s + \beta'_i(\phi))}{\Gamma(n_i^s + \alpha'_i(\phi) + \beta'_i(\phi))} \frac{\Gamma(\alpha'_i(\phi) + \beta'_i(\phi))}{\Gamma(\alpha'_i(\phi))\Gamma(\beta'_i(\phi))} \quad (4.60)$$

We now rewrite the optimization problem for empirical Bayes in 4.61 and note that for all cells, we require $n(x_i^d; \phi^d) \geq k(x_i^d; \phi^d) \geq 0$ to ensure valid distributions for $p(\theta_i|\phi, \{x_i^d\})$. We define equivalent constraints on type-specific transformations $a(x_i^d; \phi^d)$ and $b(x_i^d; \phi^d)$ in 4.62 and 4.63.

$$\phi^* = \operatorname{argmax}_{\phi^1, \phi^2, \dots, \phi^D} \prod_{i=1}^N p(y_i^s|\phi, \{x_i^d\}) \quad (4.61)$$

$$a(x_i^d; \phi^d) = \phi_1^d x_i^d + \phi_2^d \geq 0 \quad (4.62)$$

$$b(x_i^d; \boldsymbol{\phi}^d) = \phi_3^d x_i^d + \phi_4^d \geq 0 \quad \forall d \in 1, 2, \dots, D, i \in 1, 2, \dots, N \quad (4.63)$$

Finally, noting that $x_i^d \in [0, 1]$, we can rewrite the optimization in standard form as in 4.64 through 4.68. Solutions to this problem are straightforward to attain using standard solver libraries.

$$\boldsymbol{\phi}^* = \operatorname{argmax}_{\phi^1, \phi^2, \dots, \phi^D} \sum_{i=1}^N \log p(y_i^s | \boldsymbol{\phi}, \{x_i^d\}) \quad (4.64)$$

$$\text{s.t. } \phi_1^d + \phi_2^d \geq 0 \quad (4.65)$$

$$\phi_2^d \geq 0 \quad (4.66)$$

$$\phi_3^d + \phi_4^d \geq 0 \quad (4.67)$$

$$\phi_4^d \geq 0 \quad \forall d \in 1, 2, \dots, D \quad (4.68)$$

The empirical Bayes mapping described allows the flexible incorporation of multiple data types into the hierarchical beta model while preserving the ‘count’ interpretations necessary for beta-binomial conjugacy and efficient inference. Learning type-specific transformations in this fashion also allows the model to be robust to sparse and incomplete data. For example, if we only have sparse information on a feature like transformer location and thus have skewed representations of transformer distance, the presence of a nearby transformer should increase the probability of cell-level electrification; however, the absence of a nearby transformer should be less indicative of non-electrification. Empirical Bayes allows us to model this behavior through learning optimal transformations $a(x_i^d; \boldsymbol{\phi}^d)$ and $b(x_i^d; \boldsymbol{\phi}^d)$ specific to the available data.

We made an assumption early in this section regarding the treatment of cell variables θ_i as being conditionally independent from one another. While this is not true in our generative model, it may be a good approximation if local evidence is relatively strong. A more comprehensive approach would be to marginalize out the full latent structure for region and edge variables using sampling methods before applying empirical Bayes. Future work may explore the merits of taking this additional step.

Inference

Inference is performed for calculating posterior beliefs using MCMC methods. Cell variables $\boldsymbol{\theta}$ are sampled with a Gibbs sampler according to the posterior formulation in 4.69. This sampling step is highly efficient as all region, edge, and observation relationships are captured in closed-form through beta-binomial conjugacy. In addition, each θ_i is independent of $\boldsymbol{\theta}_{-i}$ conditioned on \mathbf{y}_i^C , $\boldsymbol{\mu}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\varepsilon}$, making it so that only one vectorized sampling step is required per iteration.

$$\theta_i | \mathbf{y}_i^C, \boldsymbol{\mu}, \boldsymbol{\gamma}, \boldsymbol{\varepsilon} \sim \text{Beta}(\gamma_\mu \mathbf{w}_i^\top \boldsymbol{\mu} + \gamma_\varepsilon \text{mean}_{k \in \mathcal{N}(i)} \varepsilon_k + \sum_{k=1}^{D^C} y_{ik}^C, \gamma_\mu (1 - \mathbf{w}_i^\top \boldsymbol{\mu}) + \gamma_\varepsilon (1 - \text{mean}_{k \in \mathcal{N}(i)} \varepsilon_k) + \sum_{k=1}^{D^C} (N_k^C - y_{ik}^C)) \quad (4.69)$$

Region variables $\boldsymbol{\mu}$, stickiness variables γ_μ and γ_ε , and edge variables $\boldsymbol{\varepsilon}_k$ are sampled using the random walk Metropolis Hastings algorithm. While the region and stickiness variables use the standard method described in Section 4.6.1, the edge variables are sampled using a parallelized implementation. The set of edge variables is partitioned into four groups of variables; each edge variable within a given group is independent of all others in that group conditioned on variables from the other groups. As such, variables within a given group can be updated in parallel. A representation of the groups for parallelization is shown in Fig. 4-20. Gray circles represent cell variables while diamonds represent edge variables. Edge variables are also colored according to groups of variables that can be sampled in parallel. The outline denotes a Markov blanket for the green edge variable in the center of the figure. Since no other green edge variables reside within the Markov blanket, green variables can be sampled in parallel. This applies to all other groups by symmetry.

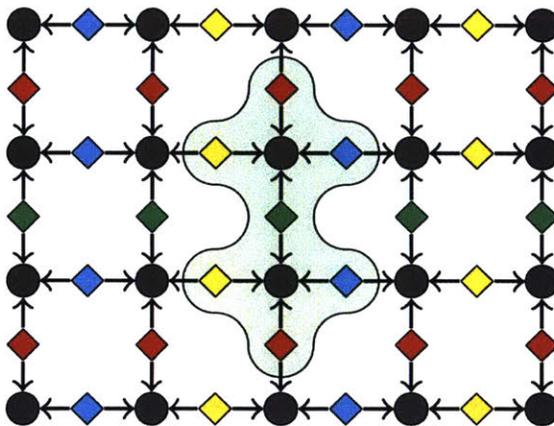
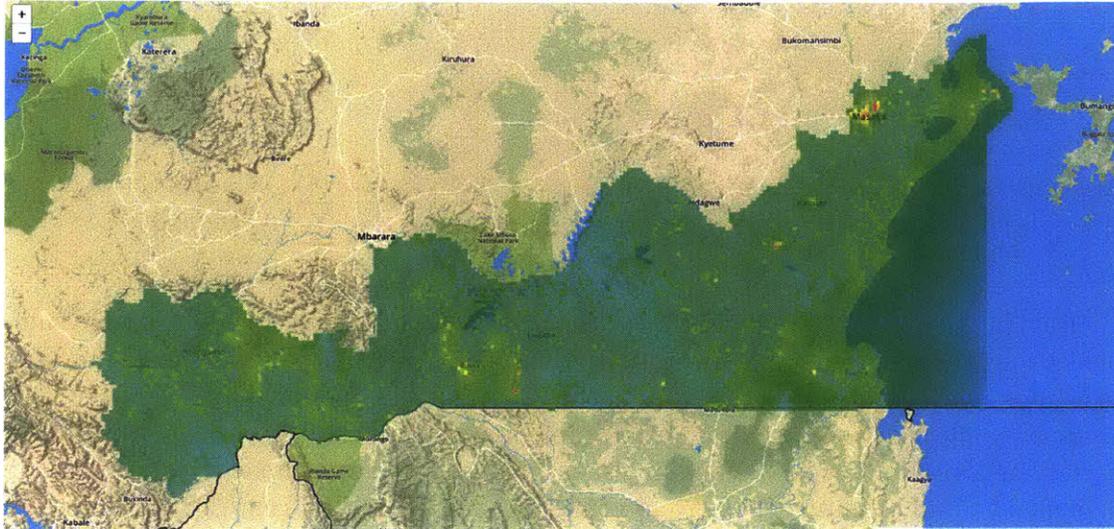
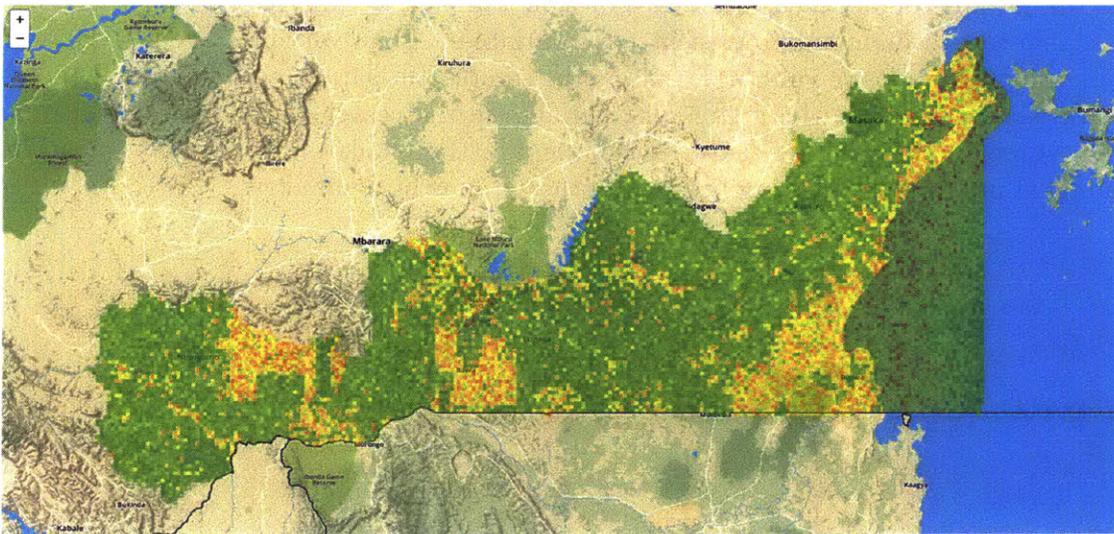


Figure 4-20: The conditional independence structure for cells enabling efficient inference via Gibbs sampling.

Using the sampling-based inference procedures described in the previous section, representations of marginal distributions and metrics for the mean, standard deviation, and entropy can be computed. Samples drawn from the posterior and statistics estimated from (e.g., quantitative measures of uncertainty such as entropy, etc.) them can also be used for a variety of infrastructure and information decision-making activities, as described in Chapter 6.



(a)



(b)

Figure 4-21: Uganda SST case study for the hierarchical beta model. (a) Mean electrification status for the Uganda SST is depicted at the cell-level. (b) Standard deviations for electrification status at the cell-level. The heat maps' color scale portrays low values in blue to high values in red.

4.6.3 Case Studies

Our implementations of case studies using the hierarchical beta model are still under development; we are specifically still implementing empirical Bayes. As such, we present preliminary results for case studies that have hand-designed data type-specific transformations in this subsection. Future work will employ the hierarchical beta model in a more finalized form.

Fig. 4-21 shows mean electrification and standard deviations for posterior estimates on cell-level electrification status for the Uganda SST. The 10,914 km² territory is divided into regular grids of length equal to 30 arc seconds (~ 1km) in accordance with the DMSP-OLS annual composite grids described in Section 4.2.1. Region-level aggregate census electrification rates are employed as in Section 4.2.3, and transformer location, transformer distance, and sparse survey data are applied at the cell-level as in Sections 4.2.2 and 4.2.2. Furthermore, cell-level building densities are computed using the FCN-derived point data from 4.2.2. Fig. 4-21 depicts means and standard deviations for electrification status at the cell-level. It is apparent that cells of high electrification probability are concentrated around major cities. In addition, cells with higher standard deviations are located in areas with fewer cell-level observations, as is the case over Lake Victoria.

Fig. 4-22 shows mean and standard deviations for posterior distributions computed from our case study in Kayonza, Rwanda. Cells are arranged in regular grids as they are with the Uganda SST case study, and one region-level aggregate electrification status metric is given for the 1,937 km² district as described in Section 4.2.4. DMSP-OLS nighttime lights are used as in the Uganda SST study, and cell-level observations including building density, transformer location, transformer distance, and cell-level electrification ground truth data are employed as presented in Sections 4.2.4, 4.2.4, and 4.2.4.

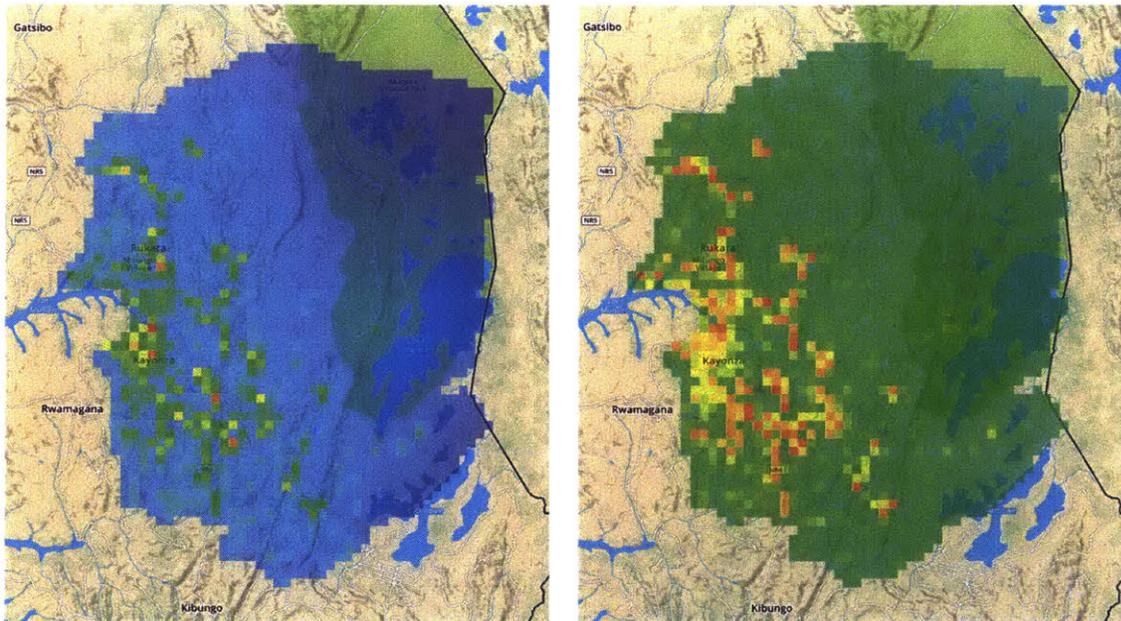


Figure 4-22: Results of applying a hierarchical beta model for electrification inference in the Kayonza district of Rwanda. Left: mean probability of electrification; Right: standard deviation. The heat maps' color scale portrays low values in blue to high values in red.

4.6.4 Discussion

Compared to the logistic regression and Gaussian process methods from Sections 4.4 and 4.5, the hierarchical beta model demonstrates greater flexibility for multiscale and multimodal inputs while maintaining efficient inference. It models spatial smoothness properties and it also features robustness to incomplete and nonconcurrent data sets. As a generative approach, the hierarchical beta furthermore enables the quantification of uncertainty. This makes it especially attractive for infrastructure and information decision-making activities, as we will expound upon in Chapter 6. One key feature pertains to the assessment of value of information (VoI) criteria, which can be used to inform strategies for rational surveying under resource constraints.

The hierarchical beta model can also be applied to other applications in sensing and information fusion where interpretability and computational efficiency are important. For energy access planning, it may be useful for the estimation of latent electricity demand. Other applications likely exist for activities related to economic development, military intelligence, agriculture, ecology, and population studies. Because the beta and binomial distributions are special cases of the Dirichlet and multinomial distributions, future work may also permit extension to multinomial observations and features of interest.

4.7 Conclusions

To the best of our knowledge, this thesis accounts for the first time probabilistic techniques have been applied to electrification status estimation with sub-village-level granularity. Previous work by Doll et al., Ellman, and Cotterman employ heuristic approaches at the building-level but make no attempt to evaluate the efficacy of their assumptions using building-level ground truth data (Doll and Pachauri, 2010b). Min et al. use logistic regression models for village-level electrification status estimation; however, the aggregated nature of their estimates are problematic for detailed planning activities (Min et al., 2013b). We argue that using probabilistic approaches for high resolution inference is essential to informing detailed electrification plans. Both mean electrification probabilities and quantified characterizations of uncertainty are necessary for representing beliefs about population characteristics over large data-scarce regions. The discriminative logistic regression model presented in this thesis falls short of providing such characterizations of uncertainty and the Gaussian process model suffers from its inability to incorporate multimodal and multiscale features.

The hierarchical beta model we present is designed to fuse a variety of inputs at an arbitrarily granular level. It can incorporate any data type that can eventually be mapped into count-data at multiple scales. It also enables the assessment of spatial correlation, which is a feature that the Gaussian process model reveals can be highly informative. We foresee other applications that can benefit from use of the hierarchical beta model, especially those pertaining to information planning activities. The model provides a framework for learning about correlations between disparate observation types and representing consistent beliefs about latent phenom-

ena. For resource-constrained activities like energy access planning in developing countries, physically attaining comprehensive measurements over large areas can be prohibitively expensive and time-consuming. Efficient probabilistic frameworks like the hierarchical beta model can be valuable for inferring desired attributes from readily available features and elucidating strategies for mitigating estimation uncertainty at minimal cost. For global challenges like the provision of universal energy access, the application of these tools can scale rapidly and drive profound social benefit.

Chapter 5

Planning and Techno-Economic Models

As outlined in Chapter 2, meaningful progress towards universal electricity access only emerges as the byproduct of complex sociotechnical systems characterized factors including political forces, evolving supply technologies, regulatory frameworks and business models, public and private finance, and consumer preferences. Even decision-makers with government and institutional-scale resources can face considerable difficulty in trying to realize desired changes.

Electricity access planning is perhaps the most salient activity that influences electricity access today; governments and distribution company planners detail plans for power systems infrastructure and direct considerable resources over time in accordance with those plans. Traditionally, this has exclusively entailed plans for extension of the central power grid; however, increasing attention is being given to off-grid technologies which may be better suited for the provision of electricity in various contexts. Planners in developing countries are now attempting to do integrated on-grid and off-grid plans.

This chapter is intended to serve as a review of modeling approaches to inform discussions of adaptive electricity access planning that will be presented in Chapter 6. This chapter is composed of three sections. The first outlines methodologies and frameworks that have been presented for electricity access planning in general. The second outlines a basic framework for thinking about techno-economic modeling for this application. Finally, the third section describes features of one such techno-economic model: the Reference Electrification Model (REM). Readers intent on learning about REM in greater depth are encouraged to review (Ellman, 2015), (Li, 2016), (Cotterman, 2017), and (Drouin, 2018).

5.1 Electrification Planning

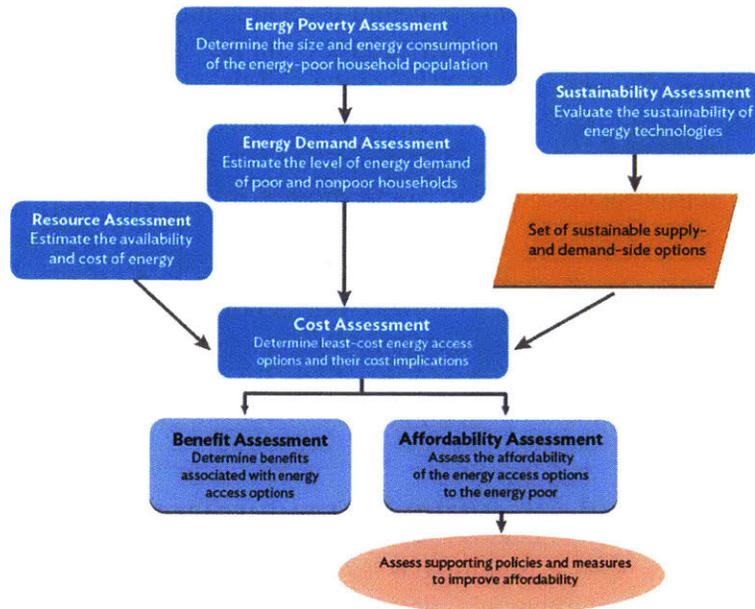
Electricity access planning (also referred to as electrification planning and electricity planning), involves assessing and designing systems for providing electricity to populations within a region of interest. Appropriate business models and system designs

are rarely obvious, given significant socio-technical complexities and the dynamic nature of service provision and demand, as discussed in Chapter 2. Such planning draws on fields within operations research, economics, management, and engineering, and in many instances can apply advanced computational methods for improved decision-making under multiple objectives.

Planning runs counter to evolutionary approaches for building power infrastructure. In the late 19th and early 20th centuries in Western society, both planning and evolutionary approaches were followed towards the proliferation of electricity infrastructure. While some networks were built according to ‘master plans’ which specify connections for major supply and demand centers, other networks evolved over time, generally in less orderly manners (Hughes, 1993). Because of the economies of scale and associated benefits of coordination that characterize power systems, planning approaches are widely seen as more favorable to evolutionary ones today.

Methods for electricity access planning are diverse, and run the gamut between ‘rational’ least-cost optimal and highly participatory approaches. Technocratic methods to rational planning benefit from their ability to scale for large regions. Computer-based techno-economic models facilitate such scaling, provided input data with detail requirements commensurate with those of the desired results. Nevertheless, these approaches fall short in their ability to incorporate context-specific considerations such as consumer traits and preferences, political goals, and institutional inertia. At the other end of the spectrum, participatory approaches are highly context-specific and aspire to involve end beneficiaries in the full lifecycle of service provision (Pritchett and Woolcock, 2004; Practical Action, 2016). While this may enable better fine-tuning relative to technocratic approaches, scaling participatory practices generally requires more time, resources, and coordination, and may be less able to take advantage technical efficiencies of greater scope. Bhattacharyya provides a detailed comparison of five different methodological options that have been used for off-grid energy supply: worksheet-based tools, optimization tools, multi-criteria decision-making tools, system-based participatory tools, and hybrid approaches. Bhattacharyya acknowledges strengths and weaknesses associated with non-hybrid methods, and ultimately recommends hybrid approaches combining two or more methods for achieving the best results (Bhattacharyya, 2012).

In addition to the methods classified above, there are a number of frameworks for electricity access planning that specify use of several of a variety of methods. For example, Shrestha et al. propose the sustainable energy access planning SEAP framework, depicted in 5-1. The SEAP framework is composed of seven assessments including an energy poverty, demand, resource, cost, benefit, sustainability, and affordability assessment (Shrestha and Acharya, 2015). Other contemporary frameworks include (Practical Action, 2016) and (Jain and Kattuman, 2015). The general purpose of these frameworks is to outline a flow for the development of more inclusive and higher quality plans.



(a)

Figure 5-1: A flow diagram depicting Shrestha et al.’s Sustainable Energy Access Planning SEAP framework. Source: (Shrestha and Acharya, 2015).

5.2 Techno-Economic Models

The importance of techno-economic modeling and modeling research for electricity access planning is arguable for a number of reasons. First of all, as alluded to in Section 5.1, techno-economic modeling does not preclude participatory measures. As such, the two paradigms can be seen as more complimentary than substitutive. For instance, optimization models can produce cost-optimal plans quickly and help to serve as a reference for planners who can subsequently deviate from them according to local contexts and input from beneficiaries. Secondly, techno-economic models can be used to rapidly compare different system designs without incurring the costs of physically building infrastructure. This is especially valuable because environments can be highly heterogeneous and computational tools may account for details that are too complicated for unassisted human planners. Third, techno-economic modeling can help decision-makers to answer regulatory and business model questions. For instance, it can provide insight for tariff structuring analyses and inform evaluations for the performance of past investments. In addition, complex models can be scaled for use by stakeholders in different parts of the world and economies of scale can be exploited with regards to model development and the compilation of parts catalogs. Finally, techno-economic models can produce clear engineering and data visualizations. These can improve stakeholder communication efforts, garner non-partisan political support, provide accountability on project execution, and help attract private investment (Howells et al., 2017).

Contemporary optimization-based techno-economic models for integrated on-grid and off-grid electrification planning include Columbia University’s Network Planner tool (Kemausuor et al., 2014), MIT and Comillas Pontifical University’s Reference Electrification Model (REM) (Ellman, 2015), and KTH Royal Institute of Technology’s Open Source Spatial Electrification Toolkit (OnSSET) (Mentis et al., 2017). Detailed comparisons between these models are given in (Cotterman, 2017) and (Drouin, 2018).

Because the work described in this thesis was developed for use in conjunction with REM and is closely affiliated with the team behind REM, aspects of subsequent discussion around techno-economic modeling and adaptive planning frameworks will be most relevant to users of the REM model. Concepts may be generalizable to other methods as well.

5.2.1 Cost-Optimal Planning Considerations

The goal of electricity access planning is to determine the most desirable ways to supply electricity to every citizen who desires it within a service area. Comprehensive plans specify a supply technology, system designs, and a business model for providing access. In most cases (and as the case for general-purpose analysis with REM), supply technologies consist of grid electrification, off-grid mini-grids, and off-grid isolated units as presented in Section 2.3. Business models are the same as those provided in Section 2.7.

Grid Electrification

Grid electrification, also referred to as grid connections and grid extensions, constitutes designs that connect individual households and buildings to the existing central grid. Complete grid extension designs may include extensions and reinforcements to low, medium, and high voltage distribution networks. They may also include models of generation assets and service reliability, specify upstream reinforcements to generation, provide detailed cost breakdowns, and enable analyses for tariff design.

Off-grid Mini-grids and Isolated Units

Plans for off-grid systems units may include designs of various levels of detail. The most salient characterization is the delineation between mini-grid and isolated systems; these calculations follow closely from considerations of off-grid economies of scale pertaining to generation assets and connection infrastructure. Other design considerations include management and tariff schemes (e.g., metering systems, business logic for implementing demand response, operations, repair and maintenance, etc.) and technical specifications (e.g., voltage and frequency controls for AC systems, etc.).

Choosing between DC and AC systems is also a topic of interest. DC systems may be able to benefit from efficiencies due to obviating the need for costly inverters that are commonly associated with AC systems. In DC systems, appliances are

more compatible with photovoltaic (PV) generation and battery technologies, which generate and store DC power. On the other hand, DC appliances are currently more costly than AC appliances, since the DC market is significantly smaller and has not yet benefited from the same economies of scale of production as AC appliances. Likewise, DC systems are also incompatible with central grids around the world. Supply infrastructure and consumer appliances would likely become redundant in the event that the central grid extends to areas serviced by DC systems.

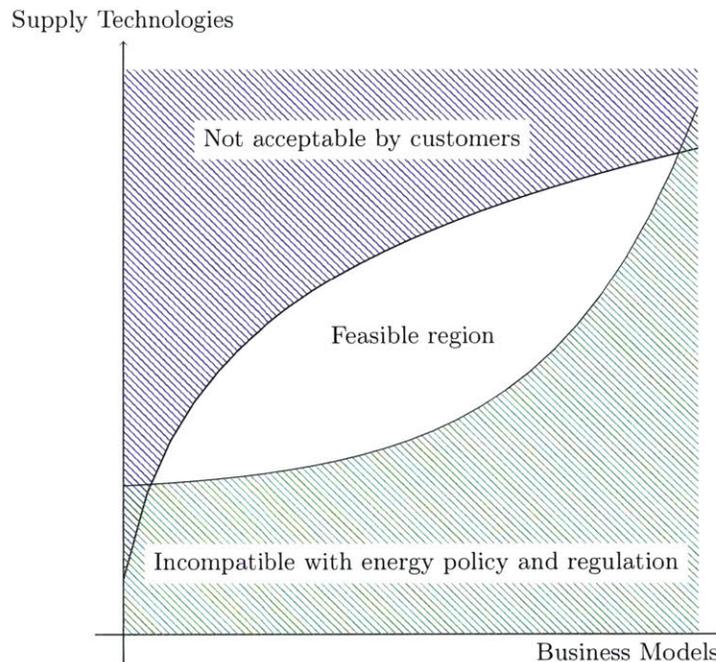


Figure 5-2: The design of techno-economic models for electrification planning may entail thinking about planning similarly to how one may think about an optimization problem. The feasible region consists of supply technologies and business models that are acceptable by customers and are compatible with appropriate policies and regulations.

The Feasible Region

The solution space for providing energy access to a given set of consumers involves all possible combinations of appropriate supply technologies and business models. Solutions may be presented at various levels of spatial granularity and plan detail. Furthermore, different solution types may be desired depending on the input data available, modeling resource availability, and the end-purpose of the plan. For instance, village-level granularities may be desired if decision-making within a country's electricity sector is very decentralized or if funding sources only require information on rough orders of magnitude for aggregate budgeting purposes. On the other hand,

highly detailed plans may be desired if decision-makers are performing comprehensive master planning or have short-term desires to break new ground on projects.

Fig. 5-2 depicts the feasible region for an illustrative range of supply technologies and business models. It is worth noting that significant numbers of constraints may limit the number of supply technologies and business model combinations that are appropriate for a given area of interest. As shown in Fig. 5-2, constraints relate to customer preferences and energy policy and regulation. For instance, consumers may be ideologically opposed to being connected to off-grid power, seeing it as inferior to grid connection. Central planners and regulators may be averse to forms of electrification that conflict with legacy policies and regulations. For example, diesel generation sets for mini-grids may be seen to be contradictory with public health and environmental objectives.

Because of the nuances with evaluating factors that may not be easily incorporated into automated methods, the selection ultimate of business models, supply technologies, and system designs may be best left to human planners. One way to judge between multiple competing plans is to use a weighted score system. A number of aspects, also known as figures of merit, may be compared. The most frequently used and easily analyzed figure of merit is likely the cost of supply; however, other considerations may include environmental, social, regulatory, participatory, and other technical factors. Weights may correspond to the level of importance of considerations, and may themselves be derived through participatory approaches.

5.3 The Reference Electrification Model

In this section, a very brief outline of REM is presented with some of its most salient characteristics to inform subsequent discussions about frameworks. As noted in the beginning of this chapter, this section is by no means exhaustive and readers should refer to (Ellman, 2015), (Li, 2016), (Cotterman, 2017), and (Drouin, 2018) for treatments of much greater depth. Though REM can be used for highly detailed system designs at small scales (referred to as ‘Local REM’), this thesis will primarily focus on regional planning which is enabled by the ‘large-scale’ use of REM.

REM uses a set of assumptions and georeferenced customer, environmental, and infrastructure data for areas of interest to determine cost-optimal supply technologies and detailed system designs for the provision of electricity access. It produces building-level local generation designs for off-grid systems (i.e., mini-grids and isolated units), clusters customers for off-grid and grid extension designs, and selects the least-cost mix of different delivery modes. With this information, a human planner can fill in other missing factors of interest, compare solutions for different scenarios in a holistic fashion, and arrive at a final electrification plan (Ellman, 2015). (Drouin, 2018) provides extensions to the core REM model for the evaluation of specific geographic and topological considerations, and (Cotterman, 2017) provides enhancements for the model’s capability to incorporate upstream network reinforcements in analyses.

Given highly accurate inputs, outputs from REM can be used as realistic project plans. Detailed technical designs can also be visualized using GIS software or web

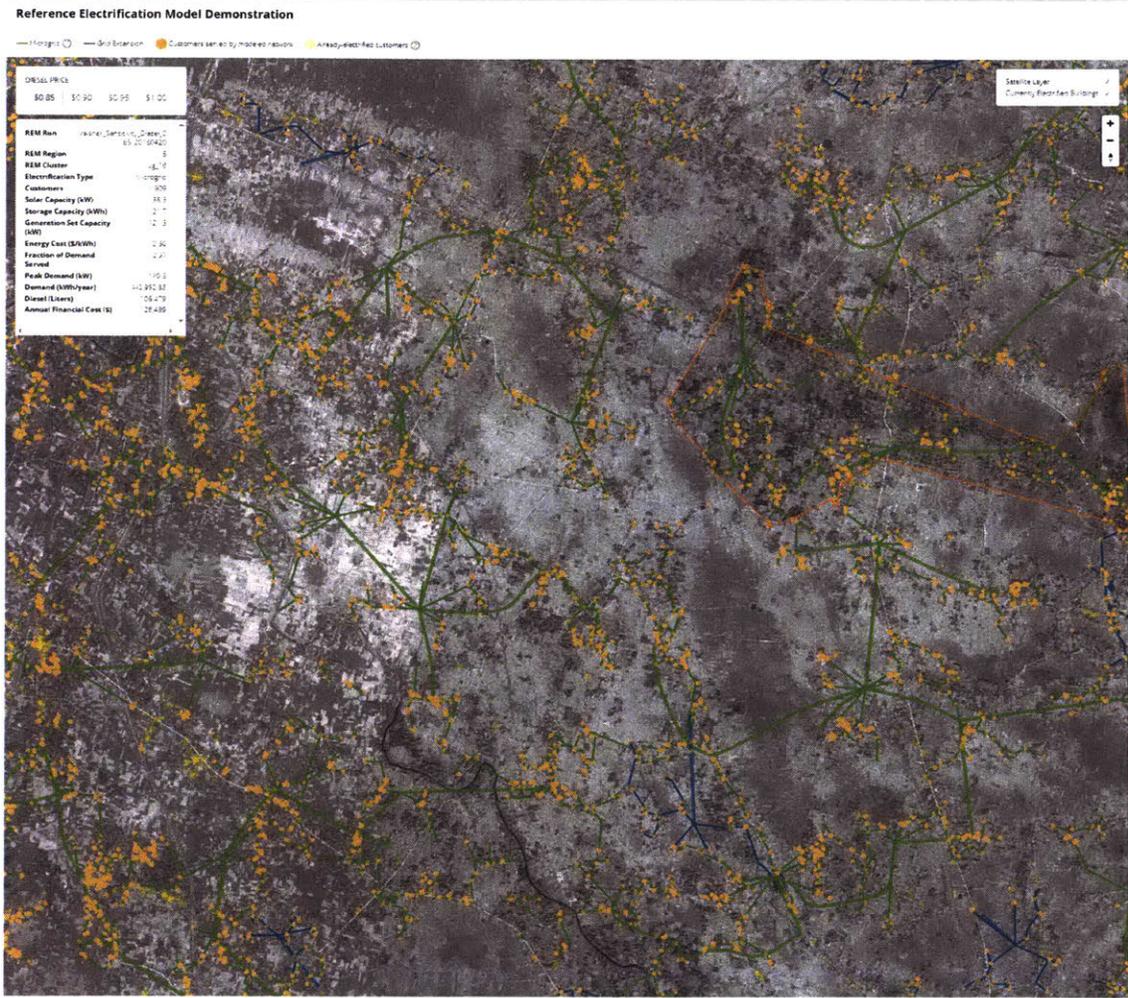


Figure 5-3: Web visualization of building localization and REM grid extension and mini-grid designs.

applications as shown in Fig. 5-3, and summary reports may be additionally compiled.

The REM model and its output is intended for use by a wide range of stakeholders. Government planners and development banks seeking to expand access through efficient infrastructure investments can use REM to help maximize the impact of scarce resources. Electricity regulators can use the output of REM viability gap benchmarking to help set appropriate tariffs and policymakers can use the tool to evaluate subsidy allocation strategies. Finally, Commercial firms can use REM to understand market sizes relevant to electrification. Companies developing generation and storage technologies can use the tool to improve product-market fits for their products and infrastructure developers can use REM to help them bid on construction tenders.

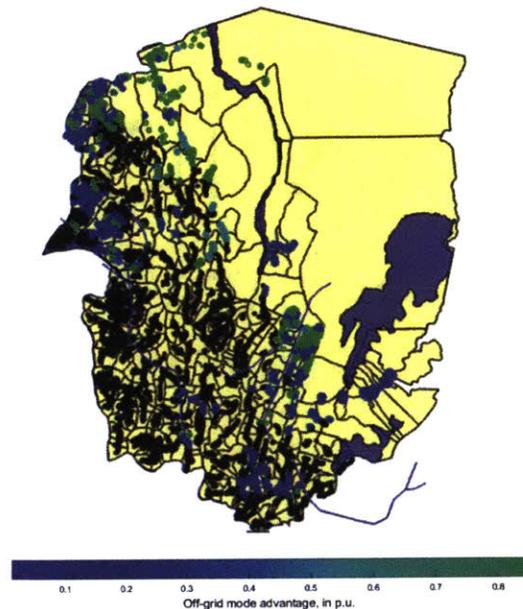


Figure 5-4: Supply technology comparisons for equivalent clusters (groups of potential customers) can give planners a relative sense for how sensitive or insensitive prescriptions are to different factors. Here, cluster points denote Off-Grid Advantage for clusters prescribed with off-grid electrification designs. Similar maps may be used to illustrate Grid Extension Advantage. Figure courtesy of Claudio Vergara. Please note that these visualizations are for demonstration purposes only and should not be interpreted as recommendations made by me, Claudio, or the MIT-Comillas Universal Energy Access Lab.

5.3.1 Off-Grid Advantage and Grid Extension Advantage

While designing networks and doing cost calculations, REM implicitly compares off-grid and grid electrification designs for given clusters of customers. After the model converges on approximations for cost-optimal systems, costs for prescribed clusters (groups of customers) may be compared to those using different electrification modes

to understand how advantageous the optimal supply technology is relative to alternatives. For instance, costs for clusters that have been designated as mini-grids and isolated systems may be compared to hypothetical grid connection costs for these same clusters. Conversely, costs for clusters prescribed with grid connections may be compared to hypothetical off-grid system costs for these clusters.

In line with comparing off-grid and grid extension costs, we propose two similar metrics that can be calculated for any clustering defined by REM: the Off-Grid Advantage and the Grid Extension Advantage, as defined by Eq. 5.1 and Eq. 5.2. For interpretability purposes, we only define the Off-Grid Advantage metric when the cost-optimal supply choice relates to off-grid systems. Likewise, we only define the Grid Extension Advantage metric when the cost-optimal supply choice is grid extension.

$$\text{Off-Grid Advantage} = \frac{\text{Cost of Grid Extension} - \text{Cost of Off-Grid}}{\text{Cost of Grid Extension}} \quad (5.1)$$

$$\text{Grid Extension Advantage} = \frac{\text{Cost of Off-Grid} - \text{Cost of Grid Extension}}{\text{Cost of Off-Grid}} \quad (5.2)$$

Fig. 5-4 depicts sample REM runs showing how Off-Grid Advantage can be visualized using color labels. Please note that these visualizations are for demonstration purposes only and should not be interpreted as recommendations made by me, Claudio, or the MIT-Comillas Universal Energy Access Lab. Grid electrification clusters are depicted by black points, while off-grid clusters are visualized as points on a blue-to-green color scale. Off-grid clusters that are more blue represent off-grid clusters that have low Off-Grid Advantage metrics, while clusters that are more green have higher Off-Grid Advantage metrics.

5.4 Conclusion

This section provides a brief overview of methodologies and frameworks for electricity access planning, presents the major benefits and limitations of techno-economic models, and outlines the general capabilities presented by the Reference Electrification Model. While well-designed and defensible plans for electricity access can only be developed when considering the full socioeconomic complexities of energy access, as described in Chapter 2, computer-based techno-economic models have the potential to still provide immense valuable to a number of stakeholders in the sector. Detailed plans, such as those generated by REM, can provide planners and others with technically robust baseline knowledge and identify high-potential projects for investment. Though providing REM with detailed input data remains a challenge for many developing countries, the building footprint extraction and electrification status estimation models presented in Chapters 3 and 4 can provide practicable approximations. Chapter 6 builds upon the background presented in this chapter and presents a framework

for iterative electricity access planning under uncertainty.

Chapter 6

Electricity Access Planning Under Uncertainty

Electricity access planning has changed in recent years with the introduction of GIS-based planning as the activity's new gold standard framework. This chapter explores how electricity access planning is poised to evolve further with the introduction of more advanced modeling and the explicit quantification of information uncertainty. Furthermore, it provides a general framework for thinking about electricity access planning under uncertainty.

Electrification plans go by a number of names in the literature. Variants of names include "Rural Electrification Master Plan," "National Electrification Master Plan," "Low Cost Rural Electrification Master Plan," "Rural Electrification Strategy and Plan," and "National Electrification Plan" (Energypedia, 2016). A recent report by the World Bank's Independent Evaluation Group (IEG) compares "classic" (also referred to as "traditional") master plans to those aided by dynamic GIS mapping activities. The IEG explains that "classic electrification master plan studies take two to three years, cost more than \$2 million to prepare, and are based on a static framework that is not readily updated" (Independent Evaluation Group, 2016). These documents can be quickly rendered out-of-date with changing technologies, policy agendas, economics for electricity access, and progress in electrification. The IEG instead recommends systematic least-cost planning supported by geographic information systems. GIS presents numerous benefits, including the ability to make more frequent updates that can describe a changing electrification landscape. Other benefits include the ability to make compelling visualizations of plans which can attract financing from national and international organizations. GIS also provides improved versatility to do comprehensive national development across various sectors and population types. Finally, managing plans using GIS have been demonstrated to dramatically save time and costs relative to classic master planning. The IEG details comprehensive geospatial planning efforts in Rwanda and Kenya, and states that they "each cost about US \$1 million and took one year to prepare" (Independent Evaluation Group, 2016). Fig. 6-1 illustrates the types of plans that GIS-based planning produced in Rwanda looking forward over a decade. Such plans have helped the Rwandese to raise \$250 MM in donor funding over 5 years (Ostojic et al., 2011).

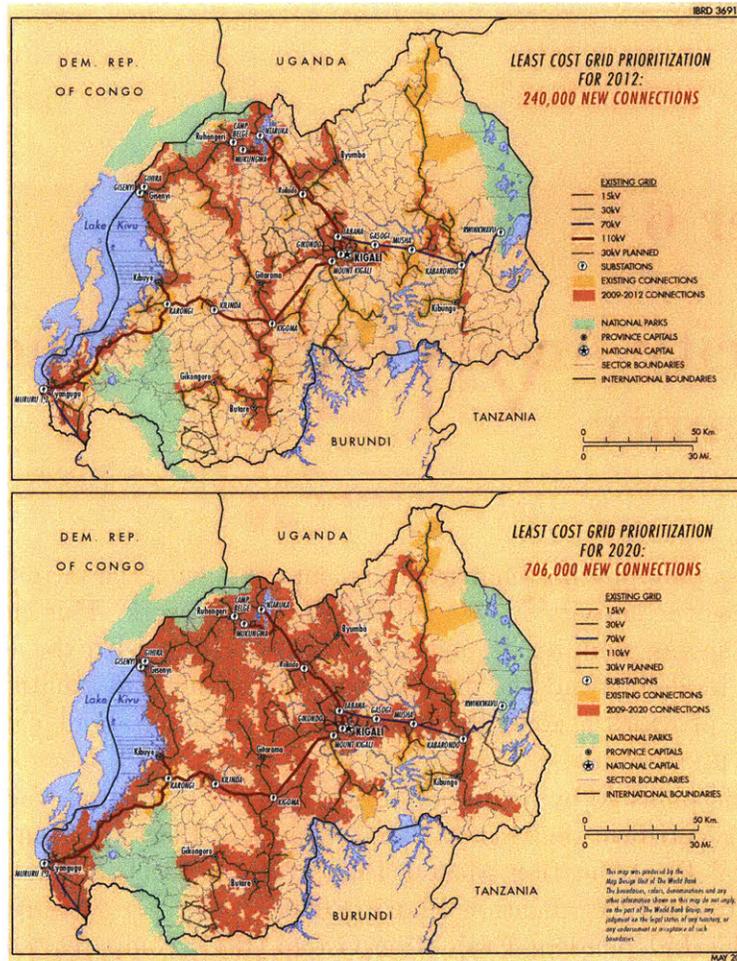


Figure 6-1: GIS-based electrification planning methodologies yielded prioritization plans in Rwanda looking out over a decade. The plans shown in this figure were developed in 2009 and prescribe areas for the rollout of electricity infrastructure out to 2012 (above) and 2020 (below). The “footprint” for grid electrification is shown in red. Off-grid systems are recommended in areas that will not receive connection via grid electrification. Figure source: (Ostojic et al., 2011).

Despite the advances that GIS-based approaches have conferred, understanding how uncertainty affects electrification planning is largely missing from the relevant literature. Uncertainty abounds in this field, especially due to the presence of data deficiencies, high resource constraints, and the dynamic nature of generation and connection technologies, populations, and institutions. The experience of the MIT-Comillas Energy Access Lab reflects that data constraints plague detailed planning efforts across sub-Saharan Africa and South Asia. Though this is changing, many planners have very little digitized geospatial information regarding electricity supply and demand in their areas of interest. Developing a large GIS database for such an area takes significant time, funding, and immediately available human capital. Taking

advantage of emerging models for probabilistic inference may help to maximize the value of limited resources and information.

We argue that even with state-of-the-art models, planning without regard for information uncertainty can stall rational planning efforts in the short-term. Planners ignorant of uncertainty characterizations may insist on comprehensive population and infrastructure surveys to be completed before funding projects, or they may decide to neglect compelling opportunities enabled by off-grid electrification. We posit that such behavior can lead to information gathering and infrastructure investment activities with lower expected return than if quantifying uncertainty is pursued simultaneously and incorporated into decision-making. Advances in computer modeling, information technology, communications (e.g., the Internet, proliferation of mobile phones, etc.), and adaptive planning will make methodologies emphasizing uncertainty quantification increasingly feasible and valuable. In the following sections, we propose and sketch the principles behind “adaptive electricity access planning,” a framework for incorporating uncertainty into planning activities for the prioritization of information and infrastructure investments. We further present a number of hypotheses supported by relevant literature:

- Economies of scale and scope can be achieved by developing frameworks and methodologies for large-scale information and infrastructure investment. The improved use of scarce resources may speed electricity provision. It may also improve planning for other social objectives such as public health and environmental concerns through explicit considerations of externalities.
- Sources of uncertainty that can be quantified reflect significant opportunities for improving adaptive computational methods through the fields of information planning and flexible design. These fields are popular in the machine learning and engineering systems communities, respectively.
- Sources of uncertainty that cannot be quantified present opportunities for improving current planning frameworks through the fields of adaptive management and planned adaptation. These fields are popular in the law and political science communities.
- Improved information can enhance stakeholder coordination, thus facilitating efforts towards the timely and efficient provision of electric power. For example, novel business models may be pursued that remove information asymmetries that currently block private investment.

6.1 Uncertainty

Although uncertainty may seem like an intuitive and easily communicable concept, there are many ways to describe it across various communities. Antunes et al. labels uncertainty as “an unintelligible expression without a straightforward description” (Antunes and Gonzalez, 2015). Nevertheless, uncertainty and its categorizations are

important to understand and communicate when making strategic decisions. In this section, we present general types of uncertainty and sources of it that affect decision-making for electricity access planning

6.1.1 Types and Sources of Uncertainty

We start by presenting different but useful ways of characterizing uncertainty, including (1) sources of uncertainty relevant to modeling and (2) analyzing knowledge and metaknowledge. We will ultimately attempt to define less common framings of uncertainty to inform a decision-making framework for adaptive electricity access planning. This will explicitly entail the definition of (3) quantifiable and unquantifiable uncertainty.

Sources of Uncertainty Relevant to Modeling

The first characterizations of uncertainty we cover are important to consider when using techniques for probabilistic inference. They describe sources of uncertainty, including inherent stochasticity, incomplete observability, and incomplete modeling. They are described below:

- **Inherent stochasticity:** Stochasticity may be inherently present in systems. Examples of this include interpretations of quantum mechanics and games of chance (e.g., hypothetical card games with card orders that are truly random, rolling die with truly random outcomes, etc.).
- **Incomplete observability:** Incomplete observations may lead to uncertainty, even if outcomes are truly deterministic. Goodfellow et al. provide the example of the Monty Hall problem for this form of uncertainty. In the problem, a game show contestant is asked to choose between three doors: two lead to goats and one to a car. Although the outcomes are deterministic based on where physical entities are located, they are effectively uncertain to the contestant.
- **Incomplete modeling:** Uncertainty can arise from imperfect models, which may discard or lose important pieces of information in observations (Goodfellow et al., 2016).

These three sources of uncertainty cover the range of those relevant to modeling. They can generally be quantified using techniques from probability and information theory and understandings of their relative magnitudes can help to inform courses of action for their amelioration. Decreasing inherent stochasticity is generally not possible without changing underlying entities and systems; however, expending resources to characterize underlying probability masses and distributions can be useful in different applications. In contrast, uncertainty that stems from incomplete observability may be mitigated through the collection of more data, and uncertainty from incomplete modeling may be reduced through the development of improved models.

Table 6.1: Johari window variant popularized by former United States Secretary of Defense Donald Rumsfeld

		Knowledge	
		<i>Knowns</i>	<i>Unknowns</i>
Metaknowledge	<i>Known</i>	<p>Known-Knowns</p> <p>(Information you have and know you have)</p>	<p>Known-Unknowns</p> <p>(Information you lack and know you lack)</p>
	<i>Unknown</i>	<p>Unknown-Knowns</p> <p>(Information you have and do not know you have)</p>	<p>Unknown-Unknowns</p> <p>(Relevant information you lack and do not know you lack)</p>

Knowledge and Metaknowledge

Another popular way to characterize uncertainty is by considering knowledge and metaknowledge, as shown in Table 6.1. In this framework, information is classified as belonging to “known-knowns,” “known-unknowns,” “unknown-knowns,” or “unknown-unknowns.” Brief descriptions are provided in the table. This representation is a variant of Luft and Ingham’s cognitive tool called the “Johari Window,” and it was popularized by former United States Secretary of Defense Donald Rumsfeld in a 2002 news briefing related to intelligence on evidence for Iraq supplying terrorists with weapons of mass destruction (Luft and Ingham, 1961; Rumsfeld, 2002). Application area aside, Rumsfeld alludes to the notion that unknown-unknowns have characteristics that are distinct from other forms of uncertainty, and thus may be particularly difficult to deal with (Rumsfeld, 2002). Known-unknowns can be addressed (and turned into known-knowns) by developing a plan and following actions to find out more information. Unknown-knowns can be dealt with through endeavors to scan for and discover missing links. Finding effective ways to address unknown-unknowns is much less obvious (Girard and Girard, 2009).

Quantifiable and Unquantifiable Uncertainty

The broad division of uncertainty types that will be used in the adaptive electricity access frameworks we present pertains to quantifiable and unquantifiable forms of uncertainty. Simple definitions of the two types of uncertainty are as follows:

- **Quantifiable Uncertainty:** Sources of uncertainty that modern methods can easily quantify in practice.

- **Unquantifiable Uncertainty:** Sources of uncertainty that modern methods cannot easily quantify in practice.

Quantifiable uncertainty may relate to inherent stochasticity, incomplete observability, and incomplete modeling, as described previously. With data points or observations, each of these types of uncertainty may be expressed quantitatively. For instance, probability theory allows for the analytical determination of the probabilities for outcomes of hypothetical card games, and measures of uncertainty such as information entropy may be computed. Similarly, models such as the hierarchical beta model, described in Section 4.6 can express measures for uncertainty from incomplete observability and incomplete modeling.

Turning to the Johari window framework presented in Table 6.1, known-unknowns can correspond to forms of both quantifiable and unquantifiable uncertainty. For instance, uncertainty around whether it will rain in a given location tomorrow can be quantified, but uncertainty regarding the existence of intelligent extraterrestrial life is arguably less measurable. On the other hand, unknown-unknowns are decidedly unquantifiable. How can one quantify something that one doesn't have information about nor even know that such information can be attained? A related concept concerns the notion of Knightian uncertainty in the study of financial markets. Knightian uncertainty pertains to information that is too vague and imprecise to be measured; it reflects limits to knowledge and the unpredictability of future events (Knight, 2012). Dizikes provides an example of Knightian uncertainty pertaining to the forecasting of airline profitability. He states, "the economic outlook for airlines 30 years from now involves so many unknown factors as to be incalculable" (Dizikes, 2010). Another related concept concerns Taleb's "theory of black swan events." The theory's name is a metaphor for major but low probability events that come as a surprise and are oftentimes unfairly rationalized by people with benefit of hindsight. Among other examples, Taleb describes the rise of the Internet, the fall of the Soviet Union, and the September 11th attacks as black swan events (Taleb, 2005, 2007).

Paul discusses the concepts of quantifiable and unquantifiable forms of uncertainty in his book, "Managing extreme financial risk: strategies and tactics for going concerns." While he acknowledges that common risk management methods can be used to protect against quantifiable uncertainty in finance, "sustainability management" must be practiced to protect against losses from unquantifiable uncertainty. Paul describes how capital is a company's primary defense against such events, and that measures must be taken for the protection and preservation of capital in response to unexpected and harmful events (Paul, 2013).

6.1.2 Sources of Uncertainty in Electricity Access Planning

Uncertainty is an inherent feature of electricity access planning, which is generally a model-assisted activity that is tightly coupled with economic development, technological development, resource constraints, human behaviors, and myriad interconnected stakeholders and institutions, as described in Chapters 2 and 5. Examples of sources

of uncertainty that affect planning are characterized as either “quantifiable” or “unquantifiable” and listed below.

- Quantifiable sources of uncertainty
 - Physical battery lifetimes
 - Battery degradation models
 - Greenhouse gas emissions from expected system operation
 - Short-term fuel prices and availability
 - Population growth
 - Electrification status
 - Electricity consumption
 - Latent electricity demand
 - Local political factors
 - Consumer perception
- Unquantifiable sources of uncertainty
 - How will technological change (e.g., supply, storage, ICT, etc.) disrupt business models 20 years from now?
 - How will resource prices (e.g., oil, natural gas, construction materials, etc.) disrupt business models 20 years from now?
 - How will socio-political factors (e.g., related to environmental displacement, religious controversies, nationalism, evolving labor markets, etc.) affect endeavors for electricity access?
 - What infectious diseases may affect populations of interest for the next 20 years and how may this affect local social fabrics?
 - Other unknown-unknowns

The “quantifiable” and “unquantifiable” characterizations will be important when thinking about appropriate methods for assisting decision-making approaches under uncertainty.

6.2 Adaptive Electricity Access Planning

In this section, we propose “adaptive electricity access planning” as a framework for electricity access planning that incorporates closed-loop processes for iterative information and infrastructure planning. Performing adaptive electricity access planning is closer to contemporary GIS planning than classical master planning (as specified in the beginning of this chapter and by the World Bank (Independent Evaluation Group, 2016)) in that it is dynamic and enables comprehensive national development

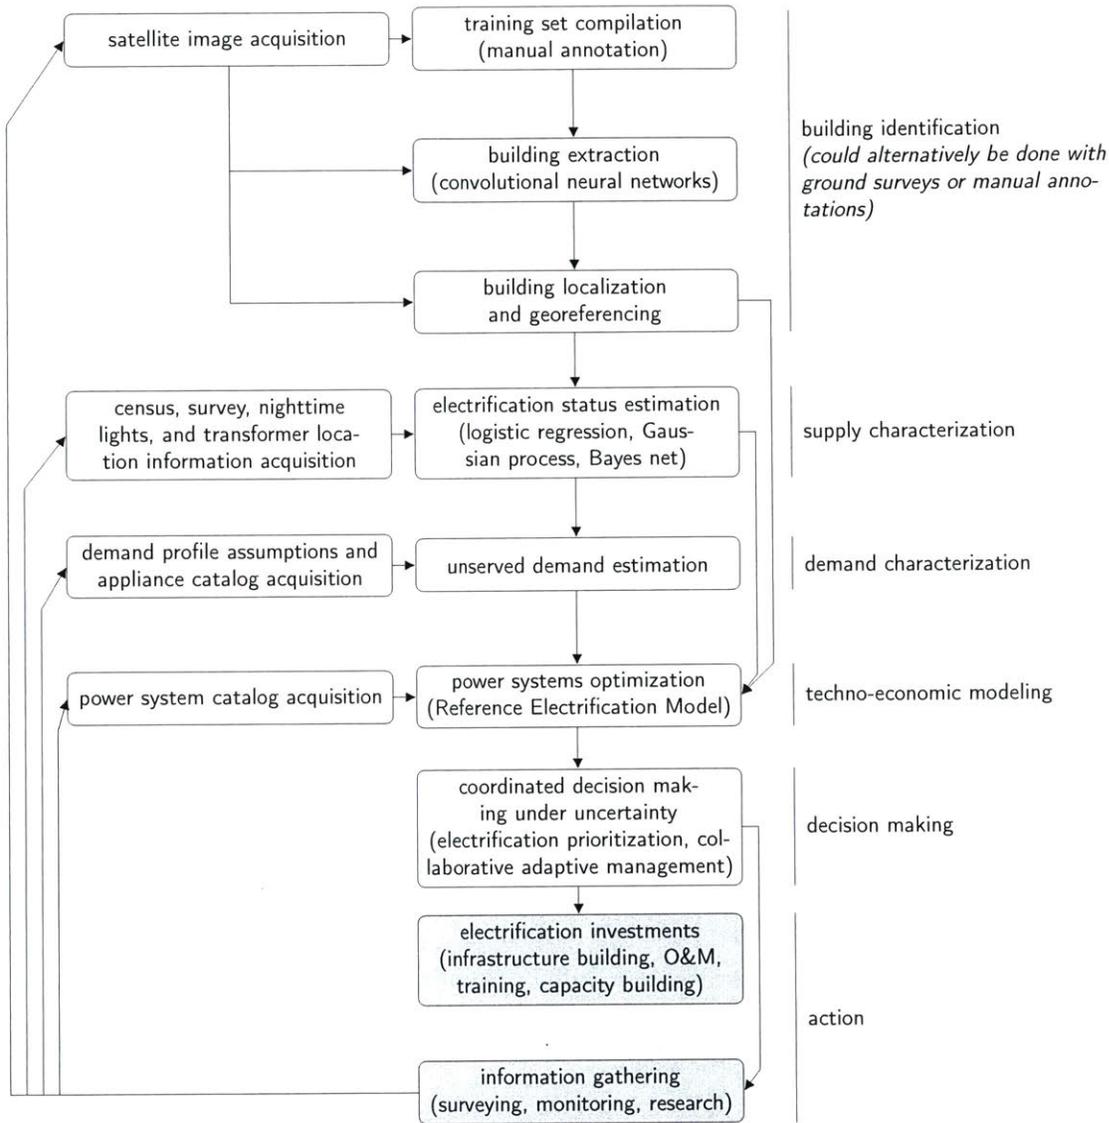


Figure 6-2: An example of a flow diagram that reflects feedback processes that are distinctive of adaptive electricity access plans. Data requirements for building identification, supply characterization, demand characterization, and techno-economic models are assessed and updated according to model-specified priorities in an “information gathering” step. “Electrification and capacity building” processes prioritize high potential impact investments given model outputs and quantifications of uncertainty. All together, the framework aims to assist planners in using the latest and greatest information sources available to continually inform regarding high priority investments and input data deficiencies.

analyses; however, it builds on contemporary approaches through explicit treatments of uncertainty. To do this, adaptive electricity access plans take advantage of probabilistic modeling methods that explicitly quantify measures of uncertainty.

Fig. 6-2 illustrates a flow diagram that outlines a representative adaptive electricity access plan. The flow diagram exemplified is composed of multiple connected processes including building identification, supply characterization, demand characterization, techno-economic modeling, and prioritization for investments in electrification and information gathering projects. While most of the connections in the diagram are relatively straightforward (e.g., building location data informs electrification status estimation, which then informs inferences on where unserved demand exists, and all of these steps ultimately provide information for techno-economic models like REM, etc.), processes for investment prioritization, coordinated decision-making, and subsequent information gathering specific to this framework may be less so. The following subsection sketches processes that planners can follow for electrification prioritization under uncertainty (corresponding to part of Fig. 6-2's "coordinated decision-making under uncertainty" process), Section 6.3 describes more detailed methods to direct continuous information gathering efforts, and Section 6.4 describes collaborate and adaptive management approaches that may improve stakeholder coordination and responsiveness to unforeseen obstacles and opportunities. Both Sections 6.3 and 6.4 provide refinements for "coordinated decision-making under uncertainty."

6.2.1 Electrification Prioritization under Uncertainty

Electrification planners usually have limited funding available at any given time for projects. As a result, they must be judicious about what types of investments to make and when. As agents representing the public interest, their top priorities are to maximize the number of people afforded access while maintaining the financial sustainability of utilities making electrification investments over time. When planning is done using tools for probabilistic inference and detailed techno-economic models, however, prioritizing electrification projects is not always obvious. Some projects may exhibit potential for electrifying consumers at low cost but correspond to areas with high model uncertainty. Other projects may have opposite characteristics. In this section, we sketch a general framework for prioritizing projects in accordance with model uncertainty. This process corresponds to "electrification prioritization," as shown in Fig. 6-2.

Naturally, high expected social benefit and recommendation certainty are desirable qualities for electrification projects. Projects that promise to yield higher expected social benefit should be prioritized over those with less compelling expectations. Likewise, all else equal, projects with higher recommendation certainty are desirable over those with lower certainty under the prevailing assumption that planners are risk averse. While it is difficult to compare the importance of expected social benefit and that of quantified recommendation certainty without detailed utility function elicitation, methodologies can be developed that approximate desired outcomes. In this section, we outline one such methodology that relies on recommendation certainty thresholds and assumptions for the propagation of uncertainty. While these assump-

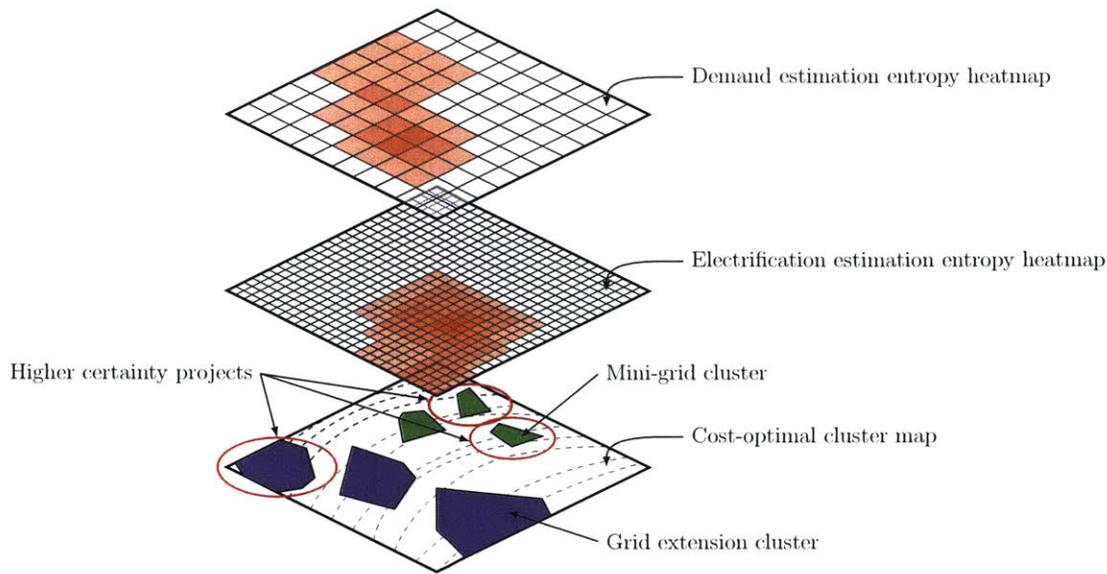


Figure 6-3: Decision-makers can use spatial position to combine multiple different layers of information including demand estimation uncertainty, electrification status estimation uncertainty, and cost-optimal infrastructure designs. Darkly shaded areas in the demand and electrification estimation entropy heatmaps denote higher cell-level uncertainty than lighter ones. ‘Higher certainty projects’ may be approximated by system designs for clusters (groupings of consumers) prescribed in areas with higher model-defined input data certainty.

tions limit the theoretical optimality of decisions, they yield benefits associated with high interpretability and ease of execution.

We describe a general framework for prioritizing investments in infrastructure combining characterizations of uncertainty and cost-optimal plans. Fig. 6-3 shows how, using spatial position, decision-makers can combine multiple different layers of information including demand estimation uncertainty, electrification status estimation uncertainty, and cost-optimal infrastructure designs. They can approximate “higher certainty projects” as designs over “clusters” (groups of consumers) that have been proposed in areas with lower quantified uncertainty from inference procedures for techno-economic model input-data. While more complicated methods of accounting for uncertainty in techno-economic models are conceivable, simply adding entropies for estimates of input data (e.g., for electrification status estimation and demand estimation, etc.) may give reasonable approximations due to the additive property of entropy under independence assumptions. Summing entropies can produce useful upper bounds on cluster-level uncertainty. A different limitation of this approach comes from assuming that techno-economic model prescriptions are spatially independent of those in neighboring regions. Common intuition related to electricity grids and economies of scale would hold that this is not true; however, if spatial dependence is considered in the inference models used (as in the hierarchical beta model presented in

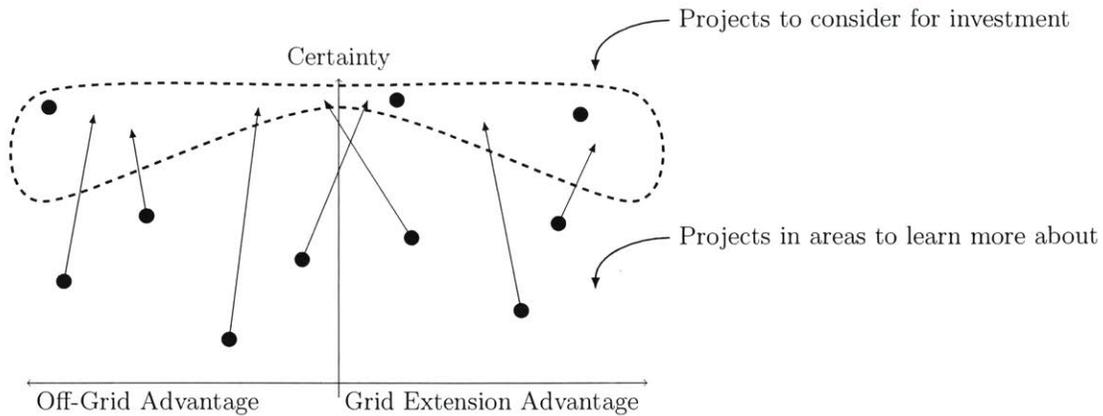


Figure 6-4: Uncertainty quantification can help to determine whether techno-economic model-defined projects should be considered for investment, or whether more information is needed before committing. Decision-makers may desire that higher certainty thresholds are used for projects with closely rivaling designs from different supply technologies.

Section 4.6), then entropies calculated for neighboring areas should be similar to one another. This indirectly mitigates issues that may arise due to spatial dependence in supply designs, since nearby clusters with coupled design decisions have similar uncertainty characteristics and will be treated alike for electrification prioritization in any case. Conversely, distant clusters will have very little affect on one another via techno-economic modeling and will have negligible spatial dependence through inference models.

Approximations for prescribed system uncertainty can be incorporated into decision-making processes in numerous ways. One straightforward method relies on using uncertainty thresholds and measures for Grid Extension Advantage and Off-Grid Advantage. As described in Section 5.3.1, prescribed system designs with high Grid Extension Advantage can be interpreted as systems for which grid extensions are significantly more cost effective than off-grid technologies. Similarly, the economics of electrification favor off-grid technologies for systems with high Off-Grid Advantage. Fig. 6-4 depicts how projects with high levels of model-quantified certainty should be considered for investment, and projects below some threshold should be designated as “projects to learn more about.” A second-order feature may be considered, concerning whether or not the system has high Off-Grid Advantage or Grid Extension Advantage scores. Systems that are only weakly specified to be a particular supply technology intuitively have higher likelihoods of changing with improved information. As such, a planner may like to have relatively higher input data certainty thresholds for investing in these projects. Fig. 6-4 also shows how additional data collection and modeling improvements can help to decrease input data uncertainty. Future work into information planning and decision-making under uncertainty can help to improve input data quality and rationalize procedures for gathering new data.

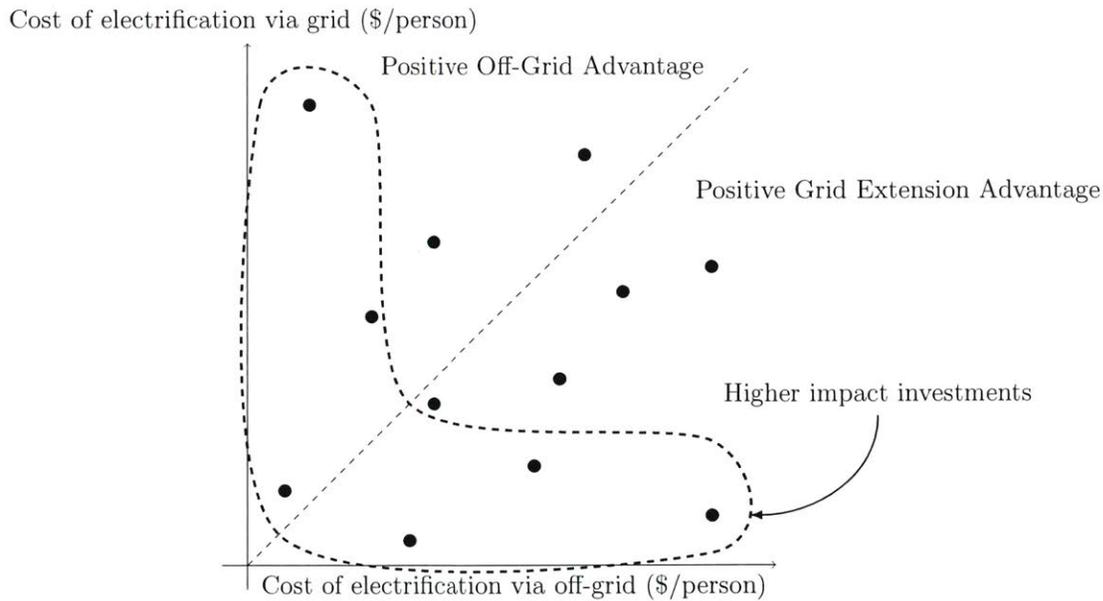


Figure 6-5: After uncertain projects are filtered, they may be screened for impact. One way to define higher impact investments is to group projects with the lowest costs of electrification on a per person basis.

Finally, Fig. 6-5 shows a final screening step for prioritizing higher impact investments over lower impact ones. This decision-making process can be visualized by plotting high certainty projects (which are the filtered output from the previous step shown in Fig. 6-4) on axes corresponding to “Cost of electrification via grid (\$/person)” and “Cost of electrification via off-grid (\$/person).” A social planner’s highest impact investments may be considered to be projects that have the lowest costs in one dimension or the other. In practice, planners can apply additional judgment to fine-tune these recommendations. This could include the incorporation of criteria that they deem most important in accordance with local values and strategic objectives.

As investments into information and infrastructure are made, progress can be conceptualized using Figs. 6-4 and 6-5. As information is improved over time, growing numbers of projects will be eligible for investment consideration. This would manifest as points migrating up the vertical axis in Fig. 6-4 and more points for investment consideration appearing on Fig. 6-5. Additionally, as the highest impact projects are completed, the cost efficiency threshold can be lowered over time to ultimately include all customer clusters and realize universal access. On Fig. 6-5, this would reflect points nearest to the horizontal and vertical axes disappearing and next highest impact projects being prioritized. Rolling out projects by cost-effectiveness can help to minimize the subsidization burden on society for a given level of electrification. Over the long-term, it can also help to maximize economic growth trajectories, as more citizens gain from the socioeconomic benefits of electrification, as described in Section 2.2. Over time, such an approach may beget virtuous cycles of improved

revenue-generation capacities and cross-subsidization for the most difficult to electrify consumers.

6.3 Dealing with Quantifiable Uncertainty

While models like the hierarchical beta model (described in Section 4.6) and frameworks for electrification prioritization under uncertainty (described in Section 6.2.1) show how model-provided quantifications of uncertainty can help to drive decision-making, techniques from the machine learning and engineering systems communities may help to achieve further benefit from quantified uncertainty. Specifically, information planning and frameworks for decision-making under uncertainty promise to help inform the most promising investment opportunities in information. Furthermore, applying methods for flexible infrastructure design may provide value by helping future stakeholders to avoid downside risks and exploit potential opportunities. In the following subsections, we provide brief overviews of avenues for future research that promise to integrate well into emerging frameworks for adaptive electricity access planning.

6.3.1 Information Planning and Decision-Making under Uncertainty

Decision-theoretic frameworks within areas of machine learning research enable a number of analyses in support of rational electricity access planning. To start, probabilistic graphical models (as described in Section 4.6.1) and utility function elicitation activities enable the computation of value of information (VoI) metrics. VoI corresponds to the increase in expected utility from making an observation or a set of observations. Performing such calculations can help agents to choose which variables in a model would likely be the most informative to observe (Kochenderfer, 2015).

Further modeling activities can help to enable the analysis of “closed-loop” electricity access planning outcomes, which stand in contrast to “open-loop” variants. Closed-loop planning employs dynamic programming algorithms to enable policies that react to the outcomes of various actions over time. Open-loop planning, on the other hand, does not account for future information and instead results in the development of a static sequence of actions. Closed-loop plans have potential benefits over open-loop plans for many real-world problems. ‘Closing the loop’ is core to the value proposition of the framework proposed in this chapter, as can be observed by comparing flow diagrams for an adaptive electricity access plan (e.g., Fig. 6-2) with those of contemporary frameworks (e.g., the SEAP framework in Fig. 5-1). Using appropriate closed-loop modeling methods promise to enhance these types of processes.

Reinforcement learning is another field that shows promise for adaptive electricity access planning. Reinforcement learning is used by agents that are performing sequential decision-making in the presence of model uncertainty, and concerns the endeavor of balancing resources spent on exploration with those spent on exploitation. Such considerations could naturally be useful to an electrification planner who

is balancing limited budgets between ground surveys and additional infrastructure projects over time. In this application, “exploration” would concern making observations to decrease model uncertainty, while “exploitation” would manifest as using whatever information is available to directly benefit customers (Kochenderfer, 2015).

While much of the literature for decision-making under uncertainty has been developed for applications in robotics, sensor placement, and other areas of engineering, these frameworks have the potential to bring significant value to electricity access planning under quantifiable uncertainty.

6.3.2 Flexible Design and Real Options

Flexible design is another field of study that has the potential to bring value to electricity access planning through the use of quantified uncertainty measures. Flexibility is defined by de Weck et al. as “the ability of a system to undergo classes of changes with relative ease. Thus, it involves a change in properties divided by resources needed to affect the change in properties.” The engineering systems literature suggests that designing systems with flexibility can help them stay viable despite uncertain futures (De Weck et al., 2011). De Neufville and Scholtes explain how flexible design can help to limit losses if negative events transpire and increase gains if positive ones do. Their book, “Flexibility in Engineering Design,” provides a four step framework that informs the development of such designs and enables the estimation of their costs and benefits (De Neufville and Scholtes, 2011).

Perhaps the most apparent application of flexible design studies pertaining to rural electrification concerns the “real option” to invest in more flexible mini-grid systems. Grid-compatible mini-grids may have higher capital costs than grid-incompatible ones; however, they also decrease the risk of stranded generation and network assets if the central grid extends to the associated service areas over the productive lifespan of the assets. Weighing the costs and benefits of these designs can help to maximize the expected value of various system designs. Flexible design considerations may analogously apply to various designs for isolated systems and the development of grid standards.

6.4 Dealing with Unquantifiable Uncertainty

Unquantifiable forms of uncertainty pose very different challenges for electricity access planning than do quantifiable forms, and as such call for different management strategies. Policy and management tools that have been proposed for dealing with such unquantifiable uncertainty, including planned adaptation, adaptive management, and Problem-Driven Iterative Adaptation.

All of these approaches share common traits regarding the use of adaptive and feedback-driven decision-making processes. Their proponents generally criticize non-adaptive practices and policies for being incapable of proactive adaptation in light of new information. McCray and Oye explain that status quo policy frameworks may be placing too much emphasis on “getting-it-right up front” and show reluctance to

reevaluate existing rules without imminent need (McCray and Oye, 2006). Such a stance leaves the public susceptible to ineffective policies for long durations of time, with decision-makers merely reacting when costly failures come to light.

Planned adaptation attempts to ameliorate this sub-optimal situation by encouraging decision-makers to appropriately recognize that uncertainties are present at the time of plan development or policy enactment. McCray et al. describe a stance based on planned adaptation to reflect, “a commitment by the decision-maker to revisit the decision at a later time in order to make any needed modifications.” Planned adaptation has also been described to draw on the concept of feedback: the acts of both sensing and controlling a process,” or having a “learning and a changing function” (McCray et al., 2010).

Adaptive management is largely similar to planned adaptation; however, the two frameworks have different vocabularies and were developed by different academic communities. Today, planned adaptation may be more commonly analyzed by researchers in the political sciences, while adaptive management may be more popular within the law and natural resources management communities. Craig and Ruhl emphasize “experimentalism” as adaptive management’s defining feature (Craig and Ruhl, 2014).

Lastly, Problem-Driven Iterative Adaptation (PDIA) is a related but newer framework aimed at promoting economic development. Though it shares many characteristics with planned adaptation and adaptive management related to feedback, experimentation, and adaptation, it emphasizes the amelioration of two system-level issues impeding development efforts: isomorphic mimicry and capability traps. Andrews et al. describe isomorphic mimicry as occurring when “governments and organizations pretend to reform by changing what policies or organizations look like rather than what they actually do.” Isomorphic mimicry is thought to weaken incentives for the pursuit of innovation and encourage rent-seeking behavior over concerns for the public good (Andrews et al., 2013). Andrews et al. go on to define capability traps as when “state capability stagnates, or even deteriorates, over long periods of time even though governments remain engaged in developmental rhetoric and continue to receive development resources.” With these debilitating issues in mind, PDIA recommends specific emphasis on local problem prioritizations and the participation of diverse stakeholder groups in the design of appropriate reforms (Andrews et al., 2013).

Though related to flexibility, adaptability is a distinct systems concept that relates to reconfiguration in response to external signals. De Weck et al. explains that adaptability “is more like the classical Darwinian concept in the sense that changes in the system are driven by changing external environments” (De Weck et al., 2011). While flexibility can be embedded into systems whose design is informed by forecasting the probabilities of future events, adaptability must be incorporated into general strategies and proactively protect against unquantifiable uncertainties. As such, this section’s main recommendation for adaptive electricity access planning is the implementation of ongoing monitoring, experimentation, and evaluation processes to ensure that electrification is proceeding in the best ways possible. For example, this could entail making institutional commitments to reassess the efficacy of regulatory, policy, and planning frameworks every 5 years. Borrowing from the PDIA and develop-

ment literatures, special emphasis should also be given to change processes that may engender isomorphic mimicry and capability traps. In this light, multi-stakeholder initiatives and local capacity building and may be especially relevant endeavors.

6.5 Platforms for Coordination

In addition to accounting for sources of uncertainty that affect planners, adaptive energy access planning may be able to beneficially integrate with software platforms that decrease uncertainty and facilitate coordination between stakeholders. Though coordination platforms may take many forms, they can potentially remove information asymmetries between public and private parties regarding electrification plans. Recent literature has found that information asymmetries can be a significant barrier to private investment in mini-grids due to uncertainty around grid extension schedules (Comello et al., 2017). Improved coordination can lower risks for developers wishing to follow unregulated business models and attract more private investment for off-grid systems.

Coordination platforms can also facilitate the exploitation of economies of scale and scope related to data procurement and prevent duplicative efforts. Activities such as infrastructure geotagging and customer surveying have fixed costs for a given area, low variable costs for scaling, and no costs for sharing. For example, implementing scaled-up versions of computer vision systems such as the one described in Chapter 3 may entail nontrivial software development and training set compilation effort, but employing the system in new areas only has costs associated with computation and satellite imagery procurement. Furthermore, both computation and satellite imagery are cheapest when performed and purchased en masse. Data collection and federation activities may be undertaken by planners and shared with other stakeholders to improve the information available for distributed electrification endeavors.

Finally, coordination platforms may constitute sensible mechanisms for the dissemination of electrification metrics and may provide accountability for improved auditing. This could go a long way towards rooting out corruption in the sector. Relatedly, different countries currently have different definitions of access, and official agencies oftentimes have perverse incentives to inflate claims regarding progress. As with any class of endeavors for improvement, electrification requires principled sensing and monitoring functions for subsequent learning and strategic action. Coordination platforms may be powerful tools to improve information integrity and protect future planning efforts.

6.6 Conclusions

This chapter wraps concepts and models introduced in previous chapters into a newly proposed framework for “adaptive electricity access planning.” Adaptive electricity access plans stand in contrast to traditional master plans and build upon GIS-based approaches by incorporating explicit processes for dealing with uncertainty for dynamic planning. Types of uncertainty are first discussed and distinctions between

quantifiable and unquantifiable sources of uncertainty are defined. By using a closed-loop process emphasizing feedback for information and infrastructure planning, it is hypothesized that more value can be made of scarce resources and progress in electrification can be accelerated. Towards this end, a specific process for incorporating cost considerations with quantified measures of uncertainty for electrification project prioritization is proposed. Next, high-level overviews are provided for how information planning and flexible design can help to make use of quantified uncertainty measures, and principles for adaptive policymaking are applied to account for unquantifiable uncertainty. Finally, a brief discussion is provided on how adaptive electricity access plans promise to enhance coordination among different stakeholders and ameliorate inefficiencies that result from information asymmetries.

Though the framework proposed in this chapter has not yet been tested, it is evident that large-scale and systems-level experimentation is required if we are to achieve universal electricity access in the next one or two decades. Status quo measures are projected to fall far short of the world's development goals by 2030 and 2040 (International Energy Agency, 2014; World Bank Independent Evaluation Group, 2015; International Energy Agency, 2017). With the advancement of technology, new opportunities are emerging that can potentially change the game. The central idea of this chapter and the unifying message of this thesis is that machine learning and detailed techno-economic models promise to be one set of technologies that should not be overlooked. Moreover, these technologies must not be developed in a vacuum. We observe that a key element that links these emerging technologies with planning processes is the explicit consideration of uncertainty. The adaptive electricity access planning framework is a first proposal for tying these elements together.

Chapter 7

Conclusions

To avoid redundantly providing overviews of the problem statement, framing, and contributions of this thesis (which was already accomplished in the introduction: Chapter 1), we conclude by emphasizing need for the codesign of regulations, policy, and planning frameworks with technological systems in order to effectively tackle complex sociotechnical problems such as universal electricity access. “Adaptive electricity access planning,” the model-driven framework for electrification planning proposed in Chapter 6, attempts to facilitate such codesign by paying special consideration to various characterizations of uncertainty. It promises to assist planners towards making the most out of limited information and resources. As energy and economic development landscapes change due to technological change, business model and regulatory innovation, and other social factors, adaptive frameworks such as the one proposed promise to provide value. Future work, significant experimentation, and careful transdisciplinary consideration is necessary to accelerate progress towards shared development goals.

Bibliography

- Aklin, M., Bayer, P., Harish, S., and Urpelainen, J. (2017). Does basic energy access generate socioeconomic benefits? A field experiment with off-grid solar power in India. *Science Advances*, 3(5):e1602153.
- Amaral, S., Câmara, G., Monteiro, A. M. V., Quintanilha, J. A., and Elvidge, C. D. (2005). Estimating population and energy consumption in Brazilian Amazonia using DMSP night-time satellite data. *Computers, Environment and Urban Systems*, 29(2):179–195.
- Amazon Web Services (2017). Amazon Mechanical Turk Getting Started Guide. <http://docs.aws.amazon.com/AWSMechTurk/latest/AWSMechanicalTurkGettingStartedGuide/Welcome.html>.
- Andrews, M., Pritchett, L., and Woolcock, M. (2013). Escaping capability traps through problem driven iterative adaptation (PDIA). *World Development*, 51:234–244.
- Antunes, R. and Gonzalez, V. (2015). A production model for construction: A theoretical framework. *Buildings*, 5(1):209–228.
- Arab, A., Hooten, M. B., and Wikle, C. K. (2008). Hierarchical spatial models. In *Encyclopedia of GIS*, pages 425–431. Springer.
- Bajcsy, R. and Tavakoli, M. (1976). Computer recognition of roads from satellite pictures. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-6(9):623–637.
- Banerjee, S., Gelfand, A. E., and Carlin, B. P. (2003). Hierarchical modeling and analysis for spatial data.
- Banerjee, S. G., Barnes, D., Singh, B., Mayer, K., and Samad, H. (2014). *Power for All: Electricity Access Challenge in India*. World Bank Publications, Washington DC.
- Bentham, J. (1823). *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press, 1907 reprinted edition.
- Bhatia, M. and Angelou, N. (2015). *Beyond Connections: Energy Access Redefined*.

- Bhattacharyya, S. C. (2012). Review of alternative methodologies for analysing off-grid electricity supply. *Renewable and Sustainable Energy Reviews*, 16(1):677–694.
- Birol, F. (2014). Energy for all will not cost the earth. In Halff, A., Sovacool, B. K., and Rozhon, J., editors, *Energy Poverty: Global Challenges and Local Solutions*, chapter 1, pages 11–20. Oxford University Press, New York, NY.
- Bischof, H., Schneider, W., and Pinz, A. J. (1992). Multispectral classification of landsat-images using neural networks. *IEEE transactions on Geoscience and Remote Sensing*, 30(3):482–490.
- Bishop, C. (2006). *Pattern Recognition and Machine Learning*.
- Bogges, J. E. (1993). Identification of roads in satellite imagery using artificial neural networks: A contextual approach.
- Burlig, F. and Preonas, L. (2016). Out of the Darkness and Into the Light? Development Effects of Rural Electrification. *Energy Institute at Haas Working Paper*, 268.
- Casella, G. (1985). An introduction to empirical Bayes data analysis. *The American Statistician*, 39(2):83–87.
- Chand, T. K., Badarinath, K., Elvidge, C., and Tuttle, B. (2009). Spatial characterization of electrical power consumption patterns over India using temporal DMSP-OLS night-time satellite data. *International Journal of Remote Sensing*, 30(3):647–661.
- Choi, M. J. and Willsky, A. S. (2007). Multiscale Gaussian graphical models and algorithms for large-scale inference. In *Statistical Signal Processing, 2007. SSP'07. IEEE/SP 14th Workshop on*, pages 229–233. IEEE.
- Columbia University Center for International Earth Science Information Network (2018). High Resolution Settlement Layer. <https://ciesin.columbia.edu/data/hrs1/>.
- Comello, S. D., Reichelstein, S. J., Sahoo, A., and Schmidt, T. S. (2017). Enabling Mini-Grid Development in Rural India. *World Development*, 93:94–107.
- Cotterman, T. L. (2017). Enhanced Techniques to Plan Rural Electrical Networks Using the Reference Electrification Model. *Massachusetts Institute of Technology*.
- Craig, R. K. and Ruhl, J. (2014). Designing administrative law for adaptive management. *Vand. L. Rev.*, 67:1.
- Cressie, N. A. (1993). *Statistics for spatial data*.
- De Neufville, R. and Scholtes, S. (2011). *Flexibility in Engineering Design*. MIT Press.

- De Weck, O. L., Roos, D., and Magee, C. L. (2011). *Engineering systems: Meeting human needs in a complex technological world*. MIT Press.
- Dean, C. L. (2015). Efficient MCMC Inference for Remote Sensing of Emission Sources. Master’s thesis, Massachusetts Institute of Technology.
- Decatur, S. (1989). Applications of Neural Networks to Terrain Classification. *Proceedings of the International Joint Conference on Neural Networks*, 1.
- DigitalGlobe (2017). GBDX. <https://www.digitalglobe.com/platforms/gbdx>.
- Dizikes, P. (2010). Explained: Knightian uncertainty. *MIT News Office*, June, 2:2010.
- Doll, C. N., Muller, J.-P., and Morley, J. G. (2006). Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics*, 57(1):75–92.
- Doll, C. N. and Pachauri, S. (2010a). Estimating rural populations without access to electricity in developing countries through night-time light satellite imagery. *Energy Policy*, 38(10):5661–5670.
- Doll, C. N. and Pachauri, S. (2010b). Estimating rural populations without access to electricity in developing countries through night-time light satellite imagery. *Energy Policy*, 38:5661–5670.
- Dollar, P., Tu, Z., and Belongie, S. (2006). Supervised learning of edges and object boundaries. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1964–1971. IEEE.
- Drouin, C. (2018). Geospatial Cost Drivers in Computer-aided Electrification Planning: The Case of Rwanda. *Massachusetts Institute of Technology*.
- Ecopia (2017). Converting Pixels Into HD Maps. <https://www.ecopiatech.com/>.
- Ellman, D. (2015). The Reference Electrification Model: A Computer Model for Planning Rural Electricity Access. *Massachusetts Institute of Technology*.
- Elvidge, C. D., Baugh, K. E., Kihn, E. A., Kroehl, H. W., Davis, E. R., and Davis, C. W. (1997). Relation between satellite observed visible-near infrared emissions, population, economic activity and electric power consumption. *International Journal of Remote Sensing*, 18(6):1373–1379.
- Energypedia (2016). List of Rural Electrification Plans.
- Forum, W. E. (2017). Partnering Against Corruption Initiative – Infrastructure and Urban Development Building Foundations for Trust and Integrity. (March):32.
- Girard, J. and Girard, J. (2009). *A leader’s guide to knowledge management: Drawing on the past to enhance future performance*. Business Expert Press.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

- Gros, A. and Tiecke, T. (2016). Connecting the world with better maps.
- Halff, A., Sovacool, B. K., and Rozhon, J. (2014). Introduction: The End of Energy Poverty. In Halff, A., Sovacool, B. K., and Rozhon, J., editors, *Energy Poverty: Global Challenges and Local Solutions*, pages 1–7. Oxford University Press, New York, NY, 1 edition.
- Haralick, R. M. (1976). Automatic remote sensor image processing. In *Digital Picture Analysis*, pages 5–63. Springer.
- Hartung, C., Lerer, A., Anokwa, Y., Tseng, C., Brunette, W., and Borriello, G. (2010). Open data kit: tools to build information services for developing regions. In *Proceedings of the 4th ACM/IEEE International Conference on Information and Communication Technologies and Development*, page 18. ACM.
- He, X., Zemel, R. S., and Ray, D. (2006). Learning and incorporating top-down cues in image segmentation. In *European Conference on Computer Vision*, pages 338–351. Springer.
- Howells, M., Rogner, H.-H., Mentis, D., and Broad, O. (2017). Energy access and electricity planning.
- Huang, C., Davis, L., and Townshend, J. (2002). An assessment of support vector machines for land cover classification. *International Journal of Remote Sensing*, 23(4):725–749.
- Huang, X. and Zhang, L. (2009). Road centreline extraction from high-resolution imagery based on multiscale structural features and support vector machines. *International Journal of Remote Sensing*, 30(8):1977–1987.
- Hughes, T. P. (1993). *Networks of Power: Electrification in Western Society, 1880-1930*. Johns Hopkins University Press.
- Idelsohn, J. (1970). A learning system for terrain recognition. *Pattern Recognition*, 2(4):293 – 301.
- Independent Evaluation Group (2016). Reliable and Affordable Off-Grid Electricity Services for the Poor: Lessons from the World Bank Group Experience. Technical report, World Bank, Washington, DC.
- Independent Evaluation Group, The World Bank Group (2015). World Bank Group Support to Energy Access, FY2000-2014: An Independent Evaluation. *The World Bank Group*, page 200.
- International Energy Agency (2014). Africa Energy Outlook: A Focus on Prospects in Sub-Saharan Africa. *World Energy Outlook Sepcial Report*, pages 1–242.
- International Energy Agency (2017). *World Energy Outlook 2017*. Organisation for Economic Co-operation and Development, OECD.

- Jain, A. and Kattuman, P. (2015). Decision-Making and Planning Framework to Improve the Deployment Success of Decentralized Rural Electrification in India. In *Sustainable Access to Energy in the Global South*, pages 129–145. Springer.
- Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., and Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794.
- Johnson, J. (2016). simple-amt. <https://github.com/jcjohnson/simple-amt>.
- Kemausuor, F., Adkins, E., Adu-Poku, I., Brew-Hammond, A., and Modi, V. (2014). Electrification planning using Network Planner tool: The case of Ghana. *Energy for Sustainable Development*, 19:92–101.
- Kendall, A., Badrinarayanan, V., and Cipolla, R. (2016). Bayesian SegNet: model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv:1511.02680v1 [cs.CV]*.
- Khandker, S. R., Barnes, D. F., and Samad, H. A. (2013). Welfare Impacts of Rural Electrification: A Panel Data Analysis from Vietnam. *Source: Economic Development and Cultural Change*, 61(3):659–692.
- Kluckner, S. and Bischof, H. (2009). Semantic classification by covariance descriptors within a randomized forest. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 665–672. IEEE.
- Kluckner, S., Mauthner, T., Roth, P. M., and Bischof, H. (2009). Semantic classification in aerial imagery by integrating appearance and height information. In *Asian Conference on Computer Vision*, pages 477–488. Springer.
- Knight, F. H. (2012). *Risk, uncertainty and profit*. Courier Corporation.
- Kochenderfer, M. J. (2015). *Decision Making Under Uncertainty: Theory and Applications*. The MIT Press, Cambridge, Massachusetts.
- Koenderink, J. J. and van Doorn, A. J. (1987). Representation of local geometry in the visual system. *Biological cybernetics*, 55(6):367–375.
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 1 edition.
- Lee, K., Brewer, E., Christiano, C., Meyo, F., Miguel, E., Podolsky, M., Rosa, J., and Wolfram, C. (2016). Electrification for “Under Grid” Households in Rural Kenya. *Development Engineering*, 1:26–35.
- Letu, H., Hara, M., Yagi, H., Naoki, K., Tana, G., Nishio, F., and Shuhei, O. (2010). Estimating energy consumption from night-time DMPS/OLS imagery after correcting for saturation effects. *International Journal of Remote Sensing*, 31(16):4443–4458.

- Li, V. (2016). The Local Reference Electrification Model: A Comprehensive Decision-Making Tool for the Design of Rural Microgrids. *Massachusetts Institute of Technology*.
- Liang-Chieh, C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. (2015). Semantic image segmentation with deep convolutional nets and fully connected crfs. In *International Conference on Learning Representations*.
- Lin, D. and Fisher, J. W. (2012). Coupling nonparametric mixtures via latent Dirichlet processes. In *Advances in Neural Information Processing Systems*, pages 55–63.
- Lindblom, C. (2018). The Science of “Muddling Through”. In *Classic Readings in Urban Planning*, pages 31–40. Routledge.
- Long, J., Shelhamer, E., and Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440.
- Luft, J. and Ingham, H. (1961). The Johari Window. *Human Relations Training News*, 5(1):6–7.
- Marshall, W. (2017). Planet to Acquire Terra Bella from Google, Sign Multi-Year Data Contract.
- Matan, O., Burges, C. J., LeCun, Y., and Denker, J. S. (1992). Multi-digit recognition using a space displacement neural network. In *Advances in Neural Information Processing Systems*, pages 488–495.
- McCray, L. E. and Oye, K. (2006). Adaptation and Anticipation: Learning from Policy Experience. In *NSF-EPA Transatlantic Uncertainty Colloquium, Washington DC*.
- McCray, L. E., Oye, K. A., and Petersen, A. C. (2010). Planned adaptation in risk regulation: An initial survey of US environmental, health, and safety regulation. *Technological Forecasting and Social Change*, 77(6):951–959.
- Mentis, D., Howells, M., Rogner, H., Korkovelos, A., Arderne, C., Zepeda, E., Siyal, S., Taliotis, C., Bazilian, M., de Roo, A., et al. (2017). Lighting the World: the first application of an open source, spatial electrification tool (OnSSET) on Sub-Saharan Africa. *Environmental Research Letters*, 12(8):085003.
- Min, B. (2015). *Power and the Vote: Elections and Electricity in the Developing World*. Cambridge University Press.
- Min, B., Gaba, K. M., Sarr, O. F., and Agalassou, A. (2013a). Detection of rural electrification in Africa using DMSP-OLS night lights imagery. *International Journal of Remote Sensing*, 34(22):8118–8141.

- Min, B., Gaba, K. M., Sarr, O. F., and Agalassou, A. (2013b). Detection of rural electrification in Africa using DMSP-OLS night lights imagery. *International Journal of Remote Sensing*, 34:8118–8141.
- Ministry of Energy and Mineral Development and Uganda Bureau of Statistics (2014). Uganda rural-urban electrification survey, 2012. (May):163.
- MIT Technology Review (2014). The Revolutionary Technique That Quietly Changed Machine Vision Forever. *MIT Technology Review*, pages 1–4. <https://www.technologyreview.com/s/530561/>.
- Mnih, V. (2013). *Machine Learning for Aerial Image Labeling*. Doctor of philosophy thesis, University of Toronto.
- Mnih, V. and Hinton, G. E. (2010). Learning to detect roads in high-resolution aerial images. In *European Conference on Computer Vision*, pages 210–223. Springer.
- Mnih, V. and Hinton, G. E. (2012). Learning to label aerial images from noisy data. In *Proceedings of the 29th International Conference on Machine Learning*, pages 567–574.
- Murphy, K. (2012). *Machine Learning: A Probabilistic Perspective*. The MIT Press, Cambridge, Massachusetts, 1 edition.
- National Geophysical Data Center (2017). Version 4 DMSP-OLS Nighttime Lights Time Series. https://ngdc.noaa.gov/eog/gcv4_readme.txt.
- National Institute of Statistics Rwanda (2014). Rwanda Fourth Population and Housing Census 2012. *Thematic Report on Population Size, Structure and Distribution*.
- Nguyen, T. T., Grabner, H., Bischof, H., and Gruber, B. (2007). On-line boosting for car detection from aerial images. In *Research, Innovation and Vision for the Future, 2007 IEEE International Conference on*, pages 87–95. IEEE.
- Nickisch, H. and Rasmussen, C. E. (2008). Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9(Oct):2035–2078.
- Noh, H., Hong, S., and Han, B. (2015). Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1520–1528.
- OpenStreetMap (2017). OpenStreetMap About. <https://www.openstreetmap.org/about>.
- Orbital Insight (2017). Orbital Insight. <https://orbitalinsight.com/>.
- Ostojic, D., Jammi, R., Barnes, D., Sanghvi, A., Mathur, S., Zhang, Y., and Durix, L. (2011). One Goal, Two Paths: Achieving Universal Access to Modern Energy in East Asia and the Pacific. Technical report, World Bank, Washington, DC.

- Paola, J. D. and Schowengerdt, R. A. (1995). A detailed comparison of backpropagation neural network and maximum-likelihood classifiers for urban land use classification. *IEEE Transactions on Geoscience and Remote Sensing*, 33(4):981–996.
- Paul, K. (2013). *Managing Extreme Financial Risk: Strategies and Tactics for Going Concerns*. Elsevier.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(Oct):2825–2830.
- Pérez-Arriaga, I. J. (2014). *Regulation of the power sector*. Springer.
- Pérez-Arriaga, I. J. (2017). New regulatory and business model approaches to achieving universal electricity access. *Papeles de Energia*, 3:37–77.
- Pérez-Arriaga, I. J., Stoner, R., Rahnema, R., Lee, S. J., Jacquot, G., González-garcía, A., Amatya, R., Brusnahan, M., and Dueñas, P. (2018). A utility approach to accelerate universal electricity access in less developed countries: A regulatory proposal. *Economics of Energy and Environmental Policy, Preprint*.
- Porway, J., Wang, K., Yao, B., and Zhu, S. C. (2008). A hierarchical and contextual model for aerial image understanding. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE.
- Practical Action (2016). Poor People’s Energy Outlook 2016: National Energy Access Planning from the Bottom Up. *Practical Action*.
- Prates, M. O., Dey, D. K., Willig, M. R., and Yan, J. (2012). Transformed Gaussian markov random fields and spatial modeling. *arXiv preprint arXiv:1205.5467*.
- Pritchett, L. and Woolcock, M. (2004). Solutions when the solution is the problem: Arraying the disarray in development. *World Development*, 32(2):191–212.
- Rasmussen, C. E. and Nickisch, H. (2010). Gaussian Processes for Machine Learning (GPML) Toolbox. *Journal of Machine Learning Research*, 11:3011–3015.
- Rasmussen, C. E. and Nickisch, H. (2016a). Documentation for GPML Matlab Co4.0. online.
- Rasmussen, C. E. and Nickisch, H. (2016b). The GPML Toolbox version 4.0. Technical Documentation.
- Rasmussen, C. E. and Williams, C. K. (2006). *Gaussian Processes for Machine Learning*, volume 1. MIT press Cambridge.
- Rawls, J. (1971). *A Theory of Justice*.
- Ray, J. (2017). WorldView 4 high-resolution satellite now open for business in space.

- Richards, J. and Jia, X. (2006). *Remote Sensing Digital Image Analysis*. Berlin, Germany, 4 edition.
- Robbie, S. (2017). Planet Launches Satellite Constellation to Image the Whole Planet Daily.
- Rumsfeld, D. (2002). DoD News Briefing–Secretary Rumsfeld and Gen. Myers. *US Department of Defense*, 12.
- Sen, A. (1992). *Inequality Reexamined*. Clarendon Press.
- Shi, K., Yu, B., Huang, Y., Hu, Y., Yin, B., Chen, Z., Chen, L., and Wu, J. (2014). Evaluating the ability of NPP-VIIRS nighttime light data to estimate the gross domestic product and the electric power consumption of China at multiple scales: A comparison with DMSP-OLS data. *Remote Sensing*, 6(2):1705–1724.
- Shrestha, R. M. and Acharya, J. S. (2015). *Sustainable Energy Access Planning*. Asian Development Bank, Manila, Philippines, 1 edition.
- Sofreco (2013). EWSA, Electricity Access Roll Out Program Combined Report. Technical report, Sofreco.
- Song, X., Fan, G., and Rao, M. (2005). Automatic CRP mapping using nonparametric machine learning approaches. *IEEE Transactions on Geoscience and Remote Sensing*, 43(4):888–897.
- Sovacool, B. K. (2014). Defining, Measuring, and Tackling Energy Poverty. *Energy poverty: Global challenges and local solutions*, 2:21–53.
- Su, H., Deng, J., and Fei-Fei, L. (2012). Crowdsourcing annotations for visual object detection. In *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence*, volume 1.
- Sustainable Energy for All (2017). Sustainable Energy for All - Our Mission. <http://www.se4all.org/>.
- Taleb, N. N. (2005). *Foiled by Randomness: The Hidden Role of Chance in Life and in the Markets*, volume 1. Random House Incorporated.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*, volume 2. Random house.
- Tiecke, T. (2016). Open population datasets and open challenges.
- Townsend, A. C. and Bruce, D. A. (2010). The use of night-time lights satellite imagery as a measure of australia’s regional electricity consumption and population distribution. *International Journal of Remote Sensing*, 31(16):4459–4480.
- Uganda Bureau of Statistics (2014). The 2014 Uganda Population and Housing Census. <http://www.ubos.org/2014-census/>.

- United Nations (2017). Sustainable Development Goal 7. <https://sustainabledevelopment.un.org/sdg7>.
- United Nations Foundation (2017). Achieving Universal Energy Access Energy. <http://www.unfoundation.com/what-we-do/issues/energy-and-climate/clean-energy-development.html>.
- Van Uum, A. (2015). DigitalGlobe Announces Availability of 30 cm Satellite Imagery to All Customers. <https://www.businesswire.com/news/home/20150225005178/en/DigitalGlobe-Announces-Availability-30-cm-Satellite-Imagery>.
- Varshney, K. R., Chen, G. H., Abelson, B., Nowocin, K., Sakhrani, V., Xu, L., and Spatocco, B. L. (2015). Targeting Villages for Rural Development Using Satellite Image Analysis. *Big Data*, 3(1):41–53.
- Wolff, E. N. (2009). *Poverty and income distribution*. Blackwell Pub.
- World Bank Independent Evaluation Group (2015). World Bank Group Support to Electricity Access, FY2000-2014: An Independent Evaluation. Technical report, World Bank, Washington, DC.
- World Energy Outlook (2017). Modern Energy for All: Why it Matters.
- Xiao, J. (2016). Professor x toolkit (profxkit).
- Xiao, J. (2017). DrawMe: A light-weight Javascript library for line drawing on a picture.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. (2010). Sun database: Large-Scale Scene Recognition from Abbey to Zoo. In *Computer Vision and Pattern Recognition*, pages 3485–3492. IEEE.
- Yu, H., Dauwels, J., Zhang, X., Xu, S., and Uy, W. I. T. (2012). Copula gaussian multiscale graphical models with application to geophysical modeling. In *Information Fusion (FUSION), 2012 15th International Conference on*, pages 1741–1748. IEEE.
- Yuan, J. (2016). Automatic building extraction in aerial scenes using convolutional networks. *arXiv preprint arXiv:1602.06564*.
- Zhang, A., Liu, X., Gros, A., and Tiede, T. (2017). Building detection from satellite images on a global scale. *arXiv preprint arXiv:1707.08952*.
- Zhao, N., Ghosh, T., and Samson, E. L. (2012). Mapping spatio-temporal changes of chinese electric power consumption using night-time imagery. *International Journal of Remote Sensing*, 33(20):6304–6320.

Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., and Torr, P. H. (2015). Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1529–1537.

Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., and Oliva, A. (2014). Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems*, pages 487–495.