



NATIONAL SECURITY RESEARCH DIVISION

CHILDREN AND FAMILIES
EDUCATION AND THE ARTS
ENERGY AND ENVIRONMENT
HEALTH AND HEALTH CARE
INFRASTRUCTURE AND
TRANSPORTATION
INTERNATIONAL AFFAIRS
LAW AND BUSINESS
NATIONAL SECURITY
POPULATION AND AGING
PUBLIC SAFETY
SCIENCE AND TECHNOLOGY
TERRORISM AND
HOMELAND SECURITY

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis.

This electronic document was made available from www.rand.org as a public service of the RAND Corporation.

Skip all front matter: [Jump to Page 1](#) ▼

Support RAND

[Purchase this document](#)

[Browse Reports & Bookstore](#)

[Make a charitable contribution](#)

For More Information

Visit RAND at www.rand.org

Explore the [RAND National Security Research Division](#)

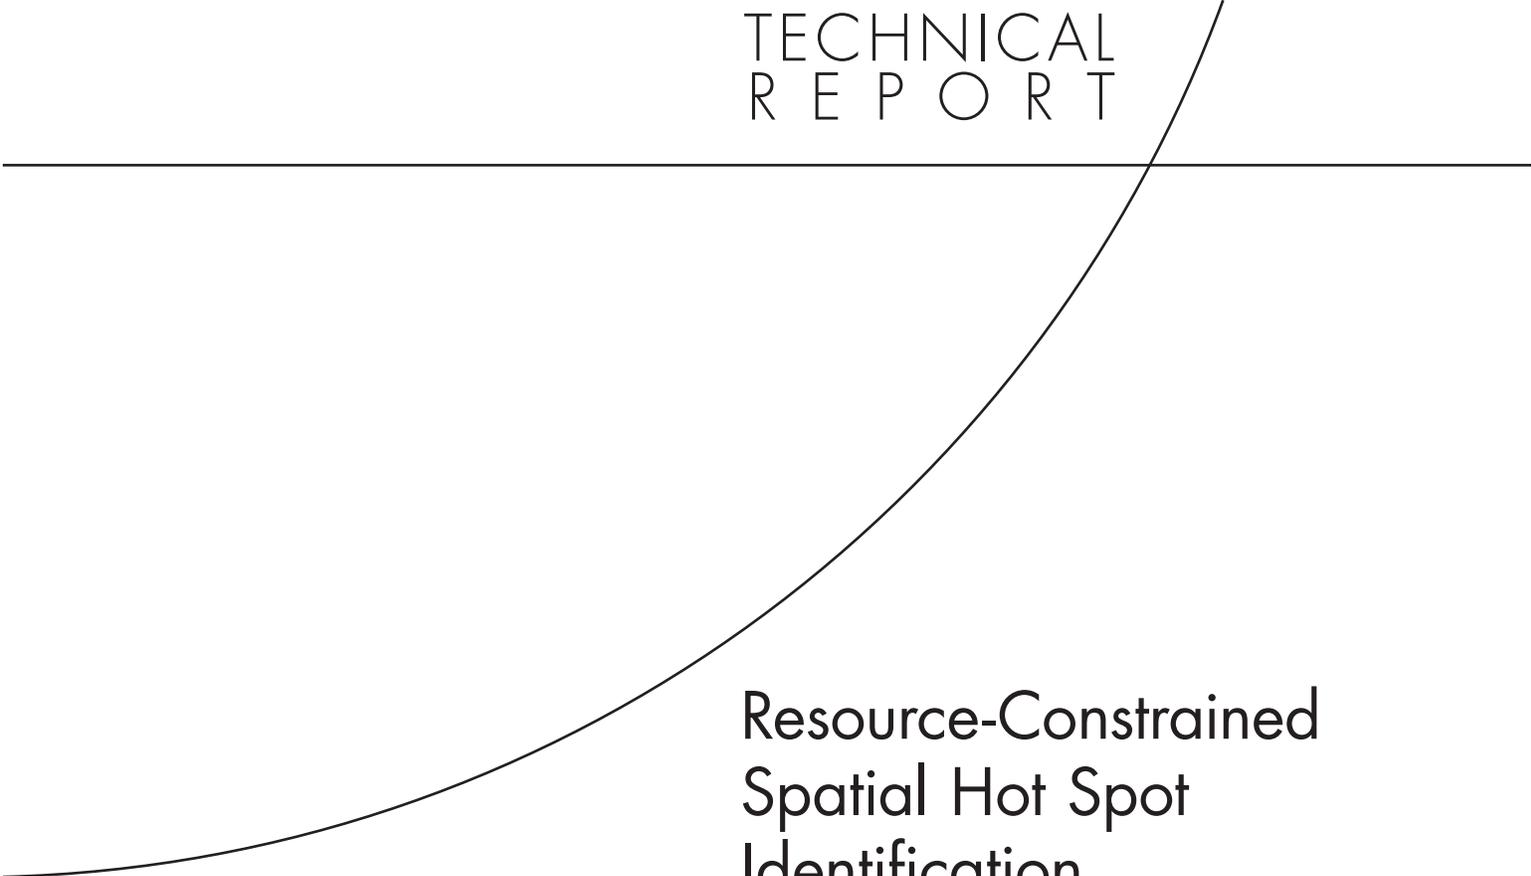
View [document details](#)

Limited Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law as indicated in a notice appearing later in this work. This electronic representation of RAND intellectual property is provided for non-commercial use only. Unauthorized posting of RAND electronic documents to a non-RAND website is prohibited. RAND electronic documents are protected under copyright law. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents for commercial use. For information on reprint and linking permissions, please see [RAND Permissions](#).

This product is part of the RAND Corporation technical report series. Reports may include research findings on a specific topic that is limited in scope; present discussions of the methodology employed in research; provide literature reviews, survey instruments, modeling exercises, guidelines for practitioners and research professionals, and supporting documentation; or deliver preliminary findings. All RAND reports undergo rigorous peer review to ensure that they meet high standards for research quality and objectivity.

TECHNICAL
R E P O R T



Resource-Constrained Spatial Hot Spot Identification

Ryan Keefe, Thomas Sullivan

Approved for public release; distribution unlimited



NATIONAL SECURITY RESEARCH DIVISION

This technical report is a product of the RAND Corporation's continuing program of self-initiated independent research. Support for such research is provided, in part, by donors and by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers.

Library of Congress Cataloging-in-Publication Data

Keefe, Ryan.

Resource-constrained spatial hot spot identification / Ryan Keefe, Thomas Sullivan.

p. cm.

Includes bibliographical references.

ISBN 978-0-8330-5137-0 (pbk. : alk. paper)

1. Cluster analysis. 2. Risk assessment—Mathematics. 3. Hazardous geographic environments—Mathematical models. 4. Piracy—Risk assessment—Aden, Gulf of—Case studies. 5. Crime forecasting—Case studies. 6. Medical geography. 7. Colon (Anatomy)—Cancer—Risk assessment—Statistical methods—Case studies. I. Sullivan, Thomas, 1965- II. Title.

QA278.K395 2011

361.1—dc22

2010050702

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2011 RAND Corporation

Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Copies may not be duplicated for commercial purposes. Unauthorized posting of RAND documents to a non-RAND website is prohibited. RAND documents are protected under copyright law. For information on reprint and linking permissions, please visit the RAND permissions page (<http://www.rand.org/publications/permissions.html>).

Published 2011 by the RAND Corporation

1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, VA 22202-5050

4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665

RAND URL: <http://www.rand.org>

To order RAND documents or to obtain additional information, contact

Distribution Services: Telephone: (310) 451-7002;

Fax: (310) 451-6915; Email: order@rand.org

Preface

A significant challenge facing policy decisionmakers tasked with combating crime, terrorism, insurgent activity, or public health risks is the scarcity of resources that can be applied to address these problems. In order to allocate limited resources, a common practice is to identify areas where the problems are more pronounced and then direct resources toward those focus areas. When the historical instances of the problem may be represented geographically, spatial analysis tools can be used to identify clusters of concentrated activity against which resources may be deployed. In the extensive body of research addressing the use of spatial analysis, the term *hot spot* has been adopted to indicate areas where there exists a greater-than-average number of historical or anticipated problem events.

In 2005, as part of the RAND Counter Improvised Explosive Device (IED) Study, the authors developed a methodology that could be used to identify IED hot spots that was constructed to match the scarce resources available to various tactical commanders in Iraq. RAND's modifications to existing spatial analysis tools allowed decisionmakers to limit the number of candidate IED hot spots and to focus on areas that conformed to the physical limits of the resources they intended to deploy against IED emplacements (e.g., sensor ranges, reachability by quick response teams during the IED emplacement stage). Additionally, the approach allowed the commanders to prioritize the reduced set of resource-constrained hot spots based on temporal patterns discerned from historical enemy IED emplacement activity.

This technical report describes a generalized version of the actionable hot spot (AHS) methodology that may find usefulness beyond the counter-IED application for which it was developed. Any decisionmaker who is faced with deploying scarce resources to geographic areas where certain types of undesirable activity or phenomena occur may find this approach useful. This approach is not intended to replace any existing spatial analysis tools but rather to augment them with the ability to conduct analysis where known constraints exist. To demonstrate the diversity of public policy

areas under which this approach may be used, this report also provides three example applications: one in domestic health care delivery (colon cancer screening in a state located in the western part of the United States), one in law enforcement (crime in a major metropolitan area), and one in the maritime domain with national security implications (piracy in the Gulf of Aden).

This technical report is a product of the RAND Corporation's continuing program of self-initiated independent research. Support for such research is provided, in part, by donors and by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers. The research was conducted within the RAND National Security Research Division (NSRD) of the RAND Corporation. NSRD conducts research and analysis on defense and national security topics for the U.S. and allied defense, foreign policy, homeland security, and intelligence communities and foundations and other non-governmental organizations that support defense and national security analysis.

For more information on the RAND National Security Research Division, see <http://www.rand.org/nsrd.html> or contact the director (contact information is provided on the web page).

Contents

Preface	iii.
Figures	vii.
Tables	ix.
Summary	xi.
Acknowledgments	xxi.
Abbreviations	xxiii.
CHAPTER ONE	
Introduction	1.
Hot Spot Definition	1.
Resource Constraints	2.
Actionable Hot Spots	3.
Mathematical Representation of the Resource Allocation Decision Problem	4.
Comparison with the Maximal Covering Location Problem	7.
Research Questions	8.
Report Organization	9.
CHAPTER TWO	
Spatially Constrained Hot Spot Identification	11.
Hierarchical Clustering	14.
Dissimilarity	15.
Creating the Distance Matrix	16.
Hierarchical Clustering with Spatial Constraints	18.
Spatial Dependence	21.
Augmenting Other Hot Spot Identification Approaches	22.
Spatially Constrained Resources with Noncircular Footprints	25.
Data Reduction	27.
Summary	28.

CHAPTER THREE

Cluster Point-Sharing	31.
The Second-Pass Approach	32.
Sub-matrix $\mathbf{A}(c)$	34.
Sub-matrixes $\mathbf{B}(c), \mathbf{B}^T(c)$	35.
Sub-matrix $\mathbf{E}(c)$	35.
Selecting Observations for Sharing	36.
Iteration over All Clusters Identified in the First Pass	37.

CHAPTER FOUR

Hot Spot Prioritization and Performance Measurement	39.
From Actionable Clusters to Actionable Hot Spots	39.
Prioritization	41.
Measuring Expected Performance	45.
Calibration	46.
Cross-Validation	47.
Summary	49.

CHAPTER FIVE

Case Studies	51.
Maritime Piracy	52.
Colon Cancer	55.
Metropolitan Crime	58.
Summary	60.

CHAPTER SIX

Implications	61.
 Bibliography	 63.

Figures

S.1. Boston Burglary Rates, 1999	xiv.
2.1. Boston Burglary Rates, 1999	12.
2.2. Unconstrained Clustering of Sample Data (single-link method) . . .	14.
2.3. Example of Single-Link Clustering	19.
2.4. Clustering Results (single-link method)	20.
2.5. Constrained Clustering Results (complete-link method)	21.
2.6. KDE Representation of a Spatial Observation	24.
2.7. Catchment Methods for Tests of Clustering	25.
2.8. Polygon Spatial Constraint Fitting	26.
2.9. Piracy Incidents in GoA/Somali Coast, 2004 – 2008	28.
2.10. Piracy Incidents in GoA/Somali Coast, July – December 2008	29.
3.1. Constrained Clustering Results (complete-link method)	32.
3.2. Final Improved Constrained Clustering	38.
4.1. Temporal Discount Functions	42.
4.2. Step-Function Temporal Prioritization	43.
4.3. Smooth Temporal Prioritization	44.
4.4. Cross-Validation and Effectiveness Measurement Time Line	49.
5.1. Piracy Incidents in the Gulf of Aden, July – December 2008	54.
5.2. KDE of Minority Colon Cancer Screening Noncompliance	57.
5.3. Daytime Burglary Events, December 16, 2008, to July 15, 2009 . . .	59.

Tables

S.1. Potential Applications of Actionable Hot Spot Methodology	xx.
1.1. Real-World Constraints	3.
2.1. Example Data	13.
3.1. Example Data Resorted for Cluster 8	33.
6.1. Potential Applications of Actionable Hot Spot Methodology	62.

Summary

Crimes, improvised explosive device (IED) attacks, disease outbreaks, and other disorder events are not spread uniformly across space or time. Maps of historical data generated by geospatial analysis often indicate localized clusters of notable events. The rich literature on the use of spatial analysis across many research fields posits several theories that attempt to explain the strength of spatial relationships among events that lead to clustering. Independent of the underlying cause of the clusters of events, a standard set of tools is available to the geospatial analyst community that enables the user to identify and interpret disorder activity. Among these toolkits, there has been increasing use of “hot spot” analysis to identify areas where clusters of local disorder events are most prominent and where appropriate resources should be deployed to deter, interrupt, or prevent further undesirable activity.

Resource-Constrained Hot Spot Identification

Hot spot analysis is frequently used to guide decisions about the deployment of resources intended to address the disorder activity. When the amount of resources is insufficient to address the entirety of the problem, hot spot analysis can also be used by decisionmakers to select areas with more pronounced problems and then allocate resources to those focus areas. However, policymakers tasked with allocating resources to address these problems often are keenly aware that the resources at their disposal have limitations that may drive the effectiveness of the various courses of action they desire to pursue. The decisionmaker who seeks to find an efficient and effective means of deploying resources to address problem areas must consider that his/her courses of actions are subject to the three types of constraints:

1. **Spatial.** The deployable asset(s) may have a fixed effective range (e.g., a visual sensor with a fan-shaped 130° field of view and 500-m range).

2. **Temporal.** The deployable asset(s) may be only deployed or effective at particular times (e.g., the visual sensor is ineffective at night).
3. **Quantity.** The number of deployable asset(s) is finite (e.g., funding exists for only two visual sensors).

In practice, since the decisionmaking consumers of standard hot spot analyses consider these types of constraints after the analysis has been completed, the assets being considered for deployment are often later determined to be an ineffective match for the hot spot. Without considering these limitations before the execution of the hot spot analysis, the resulting hot spots are often too large, inappropriately shaped, or out of synchronization with deployable resources. The term *actionable* will be used to indicate when the constrained resources are *available and appropriately matched* with the problem against which they will be deployed. This introduces a demand for “need-driven” methods that not only group data based on spatial similarity among events, but also identify actionable clusters given resource constraints (Ge et al., 2007).

Actionable Hot Spots

This research presents the *actionable hot spots*¹ (AHS) methodology. An actionable hot spot is defined as a hot spot having the same property as the standard hot spot discussed earlier (higher-than-average concentration of events in the study area), with one notable addition: *An actionable hot spot is a hot spot that has been determined to be appropriately sized, shaped, and synchronized with the cluster of disorder events against which scarce resources will be applied.* The methodology is not meant to replace existing hot spot analysis methods — rather, it is an implementable extension to existing methods that leverages standard statistical and innovative algorithms to ensure that only actionable hot spots are identified. The result of using this extension is that the decisionmaker, in addition to any exploratory spatial analysis where resources have not been applied, is presented with a list of hot spots in which his/her scarce resources can be effective. The application of constraints yields a reduced set of solutions that both are implementable and can be used to more efficiently allocate scarce resources. Naturally, before imposing constraints on hot spots that will render them actionable, an analyst may first apply a variety of standard spatial analysis tools to better understand the underlying data and their spatial distribution, and

¹This term was coined by a RAND researcher, Richard Mesic.

perhaps test some hypotheses that he/she has established to explain the reasons behind the disorder activity. After the initial exploratory analysis has been conducted and when constraints need to be introduced to guide resource decisions, the AHS approach can be used. For the decisionmaker, this represents a significant change in the way geospatial analysis is used to support their resource allocation.

Research Questions

In this report, we address three research questions:

1. Can existing geospatial tools be modified to ensure that any identified hot spots are actionable, given known spatial resource constraints?
2. Can identified actionable hot spots be prioritized so that the decisionmaker can efficiently allocate scarce resources to yield maximum effectiveness against problem areas?
3. Can the AHS methodology be applied to guide resource allocation in research areas beyond the IED application for which it was originally developed?

Hot Spot Identification

In geospatial software packages such as *CrimeStat*[®], *GeoDa*TM, and *ArcGIS*[®], the standard set of available hot spot analysis tools fall into three categories (Cameron and Leitner, 2005):²

Thematic Mapping. Concentrations of events are color-coded in discrete geographic areas that correspond to administrative boundaries (e.g., ZIP codes, Census tracts, police precincts).

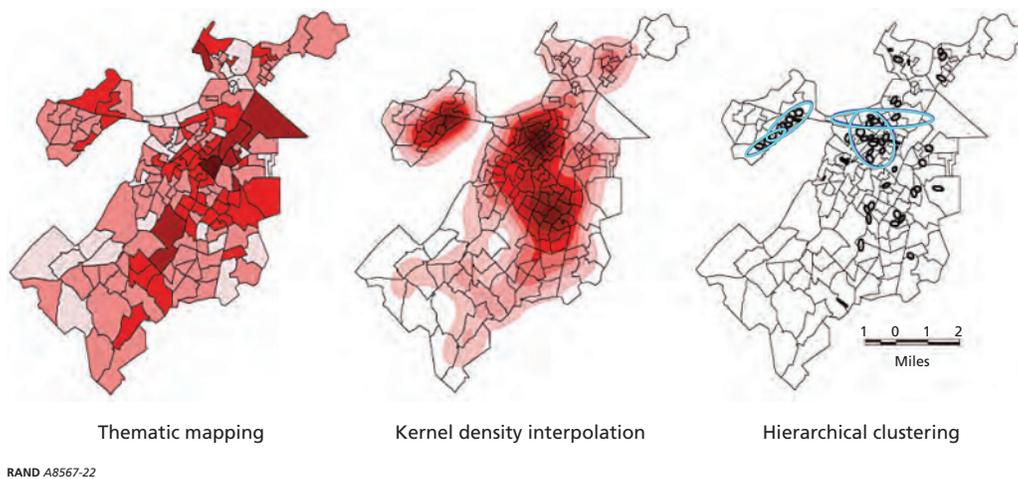
Kernel Density Interpolation. A smooth surface is overlaid on a map reflecting the concentration of actual events, and spaces between events are assigned interpolated value based on the amount of nearby events.

Hierarchical Clustering. Events are grouped according to their nearness to other events.

²This is not an exhaustive list of hot spot analysis categories, but it does include those approaches that use a sample of automated identification of hot spots rather than subjective visual interpretation.

The three categories of geospatial hot spots are illustrated in Figure S.1, reflecting maps of Boston burglary events in 1999 and provided by Cameron and Leitner (2005). The first map reflects burglary rates per 100,000 residents by Census tract, the second map reflects the density per square mile, and the final is a clustering of events contained within ellipses.³ It should be noted that the use of hierarchical clustering has been extensively discussed in spatial analysis because it is one source of the well-known *modifiable areal unit problem* (MAUP) (Openshaw, 1984) that may lead to misinterpretation of results due to the arbitrary boundaries that are used to aggregate data. Application of AHS does not resolve the MAUP, so those who interpret the results should consider that the problem may still exist.

Figure S.1
Boston Burglary Rates, 1999



Source: Cameron and Leitner, 2005.

For each of the listed categories of hot spot analysis, it is possible to modify the underlying algorithms to reflect the AHS methodology. This will result in generation of hot spots that consider the spatial, temporal, and quantity resource constraints facing the decisionmaker tasked with deploying resources to problem areas. Modification of

³Since the purpose of this illustration is to simply compare the maps resulting from the various approaches, the density and rate scales are not shown.

existing hierarchical clustering algorithms to enforce resource constraints is a rather simple exercise. Kernel density interpolation and thematic mapping approaches require considerably more effort to modify, but they, too, can be altered to consider resource constraints.

From Actionable Cluster to Actionable Hot Spots

After identifying spatially constrained clusters of disorder activity, the user will likely be left with many clusters. Three natural questions arise:

1. Which clusters are hot spots?
2. Which clusters are hotter than others?
3. Given the resource quantity constraints, against which hot spots should resources be deployed to yield maximum benefit?

To be considered an actionable hot spot, it needs to be established that the concentration of events in the clusters is greater than in other parts of the study area. The standard approach to establishing a large concentration would calculate two concentration values:

1. across the study area — the total number of events in the study area is divided by the total size of the area size (in square kilometers or miles), and the resulting concentration is denoted by c_1
2. within each cluster — the total number of events in a cluster is divided by the total size of the cluster (using the same scale that was used for the calculation of the study area) to yield a cluster concentration denoted by c_2 .

The cluster concentration is then divided by the study area concentration to yield a value $C = c_2/c_1$. If C is greater than $1.0 + \alpha$ (where $\alpha > 0$ may be defined as needed to highlight those hot spots that are distinctly different from the average density in the study area), the cluster has a higher relative concentration and is considered to be a hot spot. Actionable hot spots with higher relative concentration values are therefore considered to be “hotter” than hot spots with lower relative concentration values.

Prioritization

After actionable clusters have been determined to be actionable hot spots, the total number of these may exceed the resource quantity constraints of the decisionmaker. It is therefore required that the actionable hot spots that are candidates for resource deployment be prioritized in some fashion. Since the purpose of prioritization is to match the spatially constrained resources available in limited quantities with the problem, the prioritization should reflect the objective of the resource deployment and — if relevant to achieving that objective — temporal constraints. For example, if the objective is to reduce burglary in a small area and the deployable resource is a police patrol car available during the midnight – 8am shift, it would make little sense to put emphasis on historical events that occur during times when the patrol car is not active. A prioritization approach should put more emphasis on disorder events that occur at roughly the same time as the expected deployment of the resource.

For a given objective function and known constraints, this report proposes that each candidate actionable hot spot be weighted according to how well it is synchronized with the anticipated deployment of resources meant to combat future disorder events. The synchronization with the expected time of resource deployment can also be found through experimentation, but the basic shape of the weighting function should reflect knowledge of the deployment patterns.

Once each observation has been appropriately weighted, a cluster score may be computed, which is simply the sum of the weights in the hot spot. Prioritization then becomes simple: The actionable hot spots are ordered based on their marginal contribution to a cumulative total score (the total cumulative score will be equal to the sum of the weights for distinct events that fall within *all* identified hot spots). Resources should then be deployed first against the actionable hot spot with the highest marginal contribution to the cumulative score, followed by the one with the second highest score, etc., until the deployable resources are depleted. Since it is possible that hot spots may overlap and so events may be counted multiple times, only distinct events (those not already included in hot spots with higher marginal contributions) are counted toward marginal cluster scores. Of course, all of the events in the highest ranking hot spot will be used in the marginal score — the process of omitting nondistinct observations need only be applied to subsequent hot spots in order to accurately measure the marginal values.

Measuring Expected Performance

Although it is not possible to know how effective the resource will be once it is deployed, it is possible to use historical data to determine if the AHS-driven deployment of resources would have correctly selected areas where future events actually occurred. The performance metric is then the total number of events that occur within the recommended actionable hot spot during the resource deployment period.

For example, if the objective is to prevent burglary by sending out patrol cars to hot spots during the 8am – 4pm shift (temporal constraint), and if the cars have a patrol area of ten square city blocks (spatial constraint) and there are two patrol cars available for deployment for a period of seven days (quantity constraint), the computation of the historical metric would be done according to the following steps:

1. If the time when resource deployment will begin is represented by t (e.g., 8am on June 1, 2009), weighted actionable hot spots (given the constraints) would be computed using all relevant historical data available prior to t .
2. The two actionable hot spots with the highest weighted marginal scores would be selected for action.
3. For the next seven days beginning at time t , the number of distinct burglary events that occur within each hot spot (adjusting scores to avoid multiple counts of events that occur in more than one hot spot) during the 8am – 4pm period is counted. This is the expected performance metric.

With the performance metric, it is now possible to see whether the selection of actionable hot spots was successful. For decisionmakers comparing alternative resources for deployment, this approach will allow them to assess their potential ability to deter, disrupt, or prevent activity using various resources under consideration. Therefore, the AHS performance metric can be tested on historical data to yield an expected level of effectiveness and help choose the deployable resources that are likely to be most effective.

Case Studies

The actionable hot spot methodology was originally developed to help fight the IED problem in Iraq. Existing spatial analysis tools were modified, allowing decisionmakers to limit the number of candidate IED hot spots to areas that conformed to the physical limits of the resources tactical commanders intended to deploy

against IED emplacements. Through examples across different research areas, Chapter Five serves as a response to the third research question: *Can the actionable hot spots methodology be applied to guide resource allocation in research areas beyond the IED application for which it was originally developed?*

Any decisionmaker who is faced with deploying scarce resources to geographic areas where certain types of undesirable activity or phenomena occur may find this approach useful. This approach is not intended to replace any existing spatial analysis tools, but rather to augment them with the ability to conduct analysis where known constraints exist. To demonstrate the diversity of public policy areas under which this approach may be used, this report also provides three example applications: one in the maritime domain with national security implications (piracy in the Gulf of Aden), one in domestic health care delivery (colon cancer screening in a western U.S. state), and one in criminal justice (crime in a major metropolitan area). In each case, the actionable hot spot methodology was able to find clustering solutions that both respected the spatial, temporal, and quantity constraints and provided suggested future hot spots where events did actually occur.

We recognize that there are numerous models addressing resource allocation that have been specified for problems related to police, fire, emergency medical services, health care, etc., in addition to the IED emplacement problem. Our case studies explore research topics in which RAND is currently involved and where both the problem objectives and constraints have been clearly established by subject matter experts. Although a solution to the domestic health care delivery can be easily handled by well-known approaches such as the Maximal Covering Location Problem (MCLP) (Church and ReVelle, 1974; Church, 1984), we believe that the AHS approach provides an alternative solution that leverages commonly used hot spot identification tools and may appeal to geospatial analysts and policymakers unfamiliar with integer-programming approaches. Our approach may also add value in those types of resource allocation problems discussed in the other case studies — and perhaps additional topic areas; the prioritization phase of the AHS methodology captures shifts in spatial patterns that may occur as new target opportunities arise and/or the deployment of resources intended to interrupt future disorder events causes the actors to avoid detection. In that sense, we see the AHS approach as one possible way to address resource allocation problems when there is a repeating action-reaction exchange between those actors who deploy resources against disorder activities and those who are responsible for them.

Implications

Decisionmakers tasked with deterring, interrupting, or preventing undesired activities are limited by constraints caused by available, scarce resources; these resources often lack the ability to cover the vast geographic areas in which the problems occur. In the extensive body of research addressing the use of spatial analysis in criminal analysis, pattern recognition of insurgent and terrorist activity, and public health, the term *hot spot* has been adopted to indicate areas in which there is a greater than average number of problem events. This technical report provides a methodology that can be used to select and prioritize hot spots that can be matched with constrained resources. The methodology provides a means of measuring the expected effectiveness that would result by deploying resources against a problem using scarce resources. Not only does this approach provide a tool for aiding the decisionmaker as he/she chooses how to allocate existing resources, it also provides a mechanism for comparing the potential effectiveness of alternative resources.

The AHS methodology is not intended to replace any of the existing tools widely used by spatial analysts. Rather, it provides an enhancement to hot spot detection algorithms by enabling the geospatial analyst to match problem areas with the resources that they plan to deploy to combat the underlying problem. Users of *CrimeStat*[®], *GeoDa*TM, and *ArcGIS*[®] across many fields may find utility in this approach when they are faced with constrained resources. Originally developed for a particular application, combating IED emplacement in Iraq, the approach had obvious applications in other fields. By modifying the original application to make it generalizable across a broad array of research topics, we have created a policy decision tool that may find utility across many topical areas (see Table S.1 for a nonexhaustive list of potential applications).

Table S.1
Potential Applications of Actionable Hot Spot Methodology

Topic	Application	Deployable Resource
National security	Maritime piracy	Visual surveillance assets Armed surface ships
	Counter-IED/indirect fire	Snipers Visual surveillance assets Infrared detectors Quick reaction forces
	Insurgent network detection	Visual surveillance assets Signal direction-finding assets
Homeland security	Border integrity	Visual surveillance assets Acoustic surveillance assets Border patrol agents
Criminal justice	Policing	Police patrols Visual surveillance assets Task forces
Health	Disease prevention	Screening clinics Targeted public service campaigns
	Pandemic crises	Immunization clinics Targeted public service campaigns
Labor and population	Economic disparity	Employment programs Poverty assistance

Acknowledgments

In preparing this report, the authors sought comments and suggestions from several RAND researchers. For this, we thank Siddartha Dalal, K. Scott McMahon, Richard Mesic, Adrian Overton, Walter L. Perry, Joel Predd, Greg Ridgeway, Jess Saunders, and Henry Willis. Obtaining and reformatting the data used in the case studies would not have been possible without assistance from colleagues Seng Boey, Melissa Flournoy, Alan Fremont, Roy Gates, Jeremiah Goulka, Mark Hansen, and Oliver Wise. To increase the utility of this research, Jeffrey Sullivan generalized our existing computer code, increased its functionality by allowing the ability to overlay noncircular polygons in maps to reflect the footprints of any constrained resource, and increased the computational efficiency of many algorithms. We are thankful for the clear and useful suggestions provided by our formal reviewers, Lionel Galway from the RAND Corporation and Professor Richard L. Church from the University of California, Santa Barbara. We also acknowledge the senior leadership in RAND's National Security Research Division, who recognized and supported our efforts to generalize a RAND innovation so that it could be used to benefit a broader research portfolio: James Dobbins, Eugene Gritton, Michael Lostumbo, and Jack Riley.

Abbreviations

AHS	actionable hot spot
EO	electro-optical
GCD	Great Circle Distance
GoA	Gulf of Aden
HACM	hierarchical agglomerative clustering method
IED	improvised explosive device
KDE	kernel density estimate
MCLP	Maximal Covering Location Problem
MGRS	Military Grid Reference System
MSPA	Maritime Security Patrol Area
NNI	Nearest Neighbor Index
UAV	unmanned aerial vehicle

Introduction

The use of timely and accurate localized data to drive law enforcement operations toward more efficient and effective resource deployment is the benchmark for 21st-century policing. Strategic operations require vigilant evaluation of data through mapping technologies to identify hot spots that ultimately drive resource deployment.

— Burch and Geraci, 2009, pp. 18–20

Crimes, improvised explosive device (IED) attacks, disease outbreaks, and other disorder events are not spread uniformly across space or time. Maps of historical data generated by geospatial analysis often indicate localized clusters¹ of disorder events. The rich literature on the use of spatial analysis across many research fields posits several theories that attempt to explain the strength of spatial relationships between events that lead to the clustering of observations. Independent of the underlying cause of the clusters of disorder, a standard set of tools is available to the geospatial analyst community that enables the user to identify and interpret disorder activity. Among these toolkits, there has been increasing use of “hot spot analysis” to identify areas where clusters of local disorder events are most prominent and where appropriate resources should be deployed to deter, interrupt, or prevent further undesirable activity.

Hot Spot Definition

Research addressing the use of spatial analysis has adopted the term *hot spot* to indicate areas demonstrating a higher concentration of disorder events. In this report,

¹A *cluster* is a group of two or more data observations that are similar. In geospatial analysis, similarity — in part — reflects the spatial proximity between observations.

the formal definition of a hot spot that will be used is

- an area that contains a cluster of observations whose spatial dependence has been established using statistical testing; with a reasonable amount of confidence, it can be determined that the clustering pattern could not have occurred randomly, *and*
- the concentration of problem events in the cluster is greater than the average concentration of events in other parts of the study area.

Approaches to hot spot analysis are employed with different goals in mind (Gesler and Albert, 2000; Wilson, 2005). One approach is *general analysis*, which is used to determine if the disorder activity is clustered within the study area; the other is *focused analysis*, which is used to identify the phenomena that are clustered in a particular place in the study area. Associated with each approach is the assumption that, once the analysis has been completed and hot spots identified, it will be used to guide decisions about the deployment of resources to the areas experiencing the most problems. When the amount of resources is insufficient to address the entirety of the problem, hot spot analysis can be used by decisionmakers to select areas with more pronounced problems and then to allocate resources toward those focus areas. However, policy decisionmakers tasked with allocating resources to address these problems often are keenly aware that the resources at their disposal have limitations that may alter the effectiveness of the various courses of action they desire to pursue.

Resource Constraints

The decisionmaker who seeks to find an efficient and effective means of deploying resources to address hot spot areas must consider that his/her course of actions is subject to the three types of constraints listed in Table 1.1.

In practice, since the decisionmaking consumers of standard hot spot analyses do not consider these types of constraints in their analysis, the assets being considered for deployment are often later determined to be an ineffective match for the hot spot; the hot spots are often too large, inappropriately shaped, or out of synchronization with deployable resources. For example, if a surveillance camera with a range of 300 m is to be deployed to address disorder, hot spots with radii of 2 kilometers cannot be effective against the entire problem area using that camera. Furthermore,

Table 1.1
Real-World Constraints

Constraint	Description	Examples (ranges)
Spatial	The deployable assets have a fixed effective range — representable by a “footprint” projected onto the Earth	Surveillance camera (300-m fan-shaped footprint) Unmanned Aerial Vehicle (UAV) with an electro-optical (EO) sensor (axe-blade footprint proportional to UAV altitude) Acoustic sensor (200-m circle) Police patrol car area (10 square city blocks)
Temporal	The deployable asset may only be deployed or effective at particular times	Patrol car shifts (8 hrs) -Immunization Clinics (8am – 7pm) Thermal sensor (nighttime only)
Quantity	The number of deployable assets is finite	7 cameras 1 UAV equipped with an EO sensor

if the camera is ineffective at night, deploying it in a hot spot that reflects nocturnal disorder activity would be an inefficient use of resources. Finally, if there are only seven cameras, there is a need not only to select the hot spots where historical activity is matched with the spatial and temporal constraints but to choose the seven hot spots against which deployment of resources will be most effective. The term *actionable* will be used to indicate that constrained resources are *available and appropriately matched* with the problem against which they will be deployed. This introduces a demand for “need-driven” methods that not only group data based on spatial similarity among events, but also identify more-actionable clusters given resource constraints (Ge et al., 2007).

Actionable Hot Spots

This research presents the *actionable hot spots* (AHS) methodology. This is defined as a hot spot having the same properties as a standard hot spot discussed earlier (spatial dependence, higher-than-average concentration of events in the study area) with one notable addition: *An actionable hot spot is a hot spot of disorder activity that has been determined to be appropriately sized, shaped, and synchronized with the scarce resources that will be applied against it.*

There is a body of literature on constrained classification² addressing spatial limitations (Gordon, 1999), but since the research methods focus primarily on creating spatially contiguous clusters rather than considering resource constraints, these methods were considered irrelevant to this research. Our methodology is not meant to replace existing hot spot analysis methods — rather, it is an implementable extension of existing methods that leverages standard statistical and innovative algorithms to identify those hot spots that are actionable. The proposed approach is adaptable to handle analysis that views problem areas at various levels: specific locations, streets, neighborhoods, and large study areas. The result of using this extension is that the decisionmaker is presented with a list of hot spots in which his/her scarce resources can be effective at the appropriate level of analysis. Application of constraints yields a reduced set of solutions that both are implementable and can be used to allocate scarce resources more efficiently. Naturally, before imposing constraints on hot spots that will render them actionable, an analyst may first apply a variety of standard spatial analysis tools to better understand the underlying data,³ and their spatial distribution, and perhaps test some hypotheses that he/she has established to explain the reasons behind the disorder activity. After the initial exploratory analysis has been conducted and when constraints need to be introduced to guide resource decisions, the AHS approach can be used.

Mathematical Representation of the Resource Allocation Decision Problem

The objective of the AHS approach is to select the maximum number of expected disorder events against which constrained resources can be deployed.⁴ The resource

²When the number and identity of the classes (groupings of data) are not known in advance, the term *unsupervised classification* (Duda and Hart, 1973) is used although *clustering* is an equally acceptable term and is more often used. Similarly, although the term *constrained classification* is more widely used in the existing literature, it encompasses *constrained clustering* — the subject of this research.

³Since relationships may change for a particular type of disorder activity across the study area if the underlying environmental factors change (Haining, 2003), methods for addressing the problem must vary accordingly.

⁴The authors acknowledge Professor Richard L. Church of the University of California (Santa Barbara) for not only suggesting the need for a “clear, concise mathematical statement of the decision making problem,” but also for proposing that the problem be presented in the manner used in this report.

allocation decisionmaking problem may be stated using the following notation:

Let

- $i =$ an index representing the location of a historical disorder event
- $j =$ an index representing the center point of the location where a resource can be deployed
- $s =$ the geographic footprint within which a resource can serve, cover, or reach a location
- $w_i =$ a measure of importance of location i
- $q =$ the number of resources available to deter, disrupt, or prevent future disorder events
- $y_j =$ 1 if the resource is located at position j , 0 otherwise
- $x_i =$ 1 if the disorder activity indexed by i is within the geographic footprint of the actionable resource defined by s , 0 otherwise
- $d_{i,j} =$ 1 if the placement of the resource at position j will cause event i to be within the geographic footprint of the resource, 0 otherwise.

Note that this notation reflects some of the known constraints of the scarce resources that may be deployed against disorder events: the spatial constraint (s) and the quantity constraint (q). The constrained resource allocation decision problem is then

$$\text{Maximize } Z = \sum_i w_i x_i,$$

subject to

1. $\sum_j d_{i,j} y_j \geq x_i$ for each disorder event, i
2. $\sum_j y_j = q$
3. $y_j \in \{0,1\}$ for each j
4. $x_i \in \{0,1\}$ for each i .

It is important to recognize that temporal issues are important for some hot spot identification problems. It often is sensible to discount the importance of past events relative to very recent events; then, one can define the importance parameter, w_i , so that the importance of past events is smaller than the importance of recent events. Thus, the above model would tend to allocate resources toward areas of more-recent events and tend to ignore older events. There are also circumstances where it makes

sense to consider integrating the temporal constraints more fully in the model to reflect how well the resource deployment would be expected to be synchronized with future disorder events. We can accomplish this by modifying the above model as follows:

- $T_i = \{t \mid \text{the temporal resource scheduling period encompassing event } i\}$
- $x_{it} = 1$ if the disorder activity indexed by i is within the geographic footprint of the actionable resource defined by s at time t , 0 otherwise
- $y_{jt} = 1$ if the resource is located at position j during period t , 0 otherwise
- $d_{ijt} = 1$ if the placement of the resource at position j at time period t would cause event i to be within the geographic footprint of the resource, 0 otherwise.
- $w_{it} =$ the value of placing a resource within the geographical proximity of event i at time period t_i discounted to assign greater weight to more-recent events than to older events.

The constrained resource allocation decision problem is then

$$\text{Maximize } Z = \sum_{t \in T_i} \sum_i w_{it} x_{it}$$

subject to

1. $\sum_j d_{ijt} y_{jt} \geq x_{it}$ for each disorder event i and scheduling period t
2. $\sum_t \sum_j y_{jt} = q$
3. $y_{jt} \in \{0,1\}$ for each j
4. $\sum_t x_{it} \leq 0$ for each i .

The above model considers both geographic and temporal proximity when making the resource allocation decision (Church, personal correspondence, April 12, 2010). Note that the last constraint ensures that an event is counted only once during the identification of actionable hot spots.⁵ Most important, this model reflects all the known constraints of the scarce resources that may be deployed against disorder events: the spatial constraint, the quantity constraint, and the temporal constraint.

⁵There are cases in which a resource can be used to provide some deterrence against future events. In such cases, a more flexible, composite model might be introduced which would allocate resources based on their value toward reducing both historical and future events. This is an area in which we intend to conduct additional research in the future.

An example of how this problem might be applied would be a resource allocation problem associated with maritime piracy (this example is further developed and explained in this report’s “Case Studies” section found in Chapter Five). In that example, the objective is to locate a naval destroyer at a point in the Gulf of Aden at time t , y_{jt} , where it is within 20 nautical miles of expected future piracy activities and therefore able to deter, disrupt, or prevent those activities from occurring. The goal of the resource allocation problem is then to select a position, y_{jt} , where piracy events have occurred at greater intensity than other areas in the Gulf of Aden (as measured by $\sum_{t \in T_i} \sum_i w_{it} x_{it}$, which encompasses a measure of importance for each observation x_{it} — reflecting the degree of synchronization between the expected time when the naval destroyer may be deployed and the time when the future piracy events might occur, with perhaps more emphasis put on areas where piracy events have occurred more recently) and where those historical events are within the footprint of the deployed resource. Note that, if historical events are discounted based on the time that has elapsed since they occurred, more importance is given to recent events. The implicit assumption is that areas where events have occurred recently in higher concentrations than in the overall study area correspond to areas where additional events are expected to occur during the resource deployment period.

Comparison with the Maximal Covering Location Problem

The above mathematical model, with the exception of consideration of the temporal relevance portion, t , of the parameter w_{it} , is the well-known Maximal Covering Location Problem (MCLP), defined by Church and ReVelle (1974), extended by Church (1984), and originally used to allocate fire stations to maximize coverage of demand area. It has been applied in a variety of research areas — criminal analysis, health care delivery, advertising, emergency services, and biological reserve design — and off-the-shelf software implementations of the MCLP are readily available (Church, personal correspondence, December 6, 2009).

As evidenced by the ease in which the AHS problem can be represented using the MCLP problem formulation, the similarities between the two approaches are obvious. The AHS approach does subtly differ from the MCLP in that it aims not only to define locations for deployment of resources but to dynamically relocate them as the underlying disorder activity changes in intensity or location. Locations of disorder events, such as IED emplacement, crime, and maritime piracy, tend not be stationary but rather to shift as new target opportunities arise and/or the deployment of resources intended to interrupt disorder events causes the actors to avoid detection.

By allowing the relative importance of disorder events to lessen over time (a more formal method for defining the relative importance of the event, w_{it} , will be discussed in Chapter Four), the AHS version of the covering problem can also adapt to dynamic spatial shifts in coverage demand and indicate areas where events are expected to occur in the near future.

As the use of hot spot identification tools becomes more widely used to inform resource allocation decisions, the hot spots identified using these common tools can be the basis for assigning importance values to the MCLP. In that sense, AHS can be seen not as an alternative approach to the MCLP, but rather as a way in which hot spot identification tools can be leveraged to populate the parameter values used by the MCLP. The AHS approach differs from that of the MCLP, not only in the way in which coverage demand is determined based on the relative importance of disorder events but also in the algorithms used. As will be seen in later chapters, the AHS approach employs a simple, common hierarchical clustering method to identify hot spots while the MCLP uses a simple integer-programming algorithm. It is unclear at this time what a comparison of the relative performance of the two approaches would yield for case studies reflecting common types of disorder events, but such a comparison is one that we plan to undertake in the future. Similarly, we intend to revisit the IED emplacement problem that motivated this research to determine how well the MCLP approach performs and to better understand whether the counter-IED user community would be amenable to using such an approach.

Research Questions

In this report, we address three research questions:

1. Can existing geospatial tools be modified to ensure that any identified hot spots are actionable, given known spatial resource constraints?
2. Can identified actionable hot spots be prioritized so that the decisionmaker can efficiently allocate scarce resources to yield maximum effectiveness against problem areas?
3. Can the AHS methodology be applied to guide resource allocation in research areas beyond the IED application for which it was originally developed?

Report Organization

The remainder of this report is organized as follows. Chapter Two describes the methodology used to identify hot spots subject to spatial resource constraints and compares it with existing methods used by the geospatial analyst community. The chapter also discusses methods for reducing the size of the data set to reduce visual clutter and computational complexity in order to facilitate identification of hot spots. Chapter Three presents a method in which the clustering solutions generated by standard algorithms may be improved to include observations that were excluded from hot spots in an effort to increase computational efficiency. Chapter Four presents a method by which historical and predicted incidents may be weighted in order to prioritize the set of “actionable hot spots” to yield maximum effectiveness and describes an approach to calibrating model parameters that better match resources with the underlying problem being addressed and provides a means by which performance of the AHS approach can be measured. Chapter Five presents three case studies that span a broad range of research topics: maritime piracy, health care delivery, and law enforcement. Chapter Six discusses the implications of this research and proposes additional areas in which it may be applied.

Spatially Constrained Hot Spot Identification

As defined earlier, a hot spot is a special type of cluster where spatial dependence has been established and a higher than average concentration of activity has occurred. The standard set of hot spot analysis tools available in such geospatial software packages as *CrimeStat*[®], *GeoData*TM, and *ArcGIS*[®] (Cameron and Leitner, 2005) fall into three categories:¹

Thematic Mapping. Concentrations of events are color-coded in discrete geographic areas that correspond to administrative boundaries (e.g., ZIP codes, Census tracts, police precincts).

Kernel Density Interpolation. A smooth surface is overlaid on a map reflecting the concentration of actual events, and spaces between events are assigned interpolated value based on the amount of nearby events.

Hierarchical Clustering. Events are grouped according to their nearness to other events.

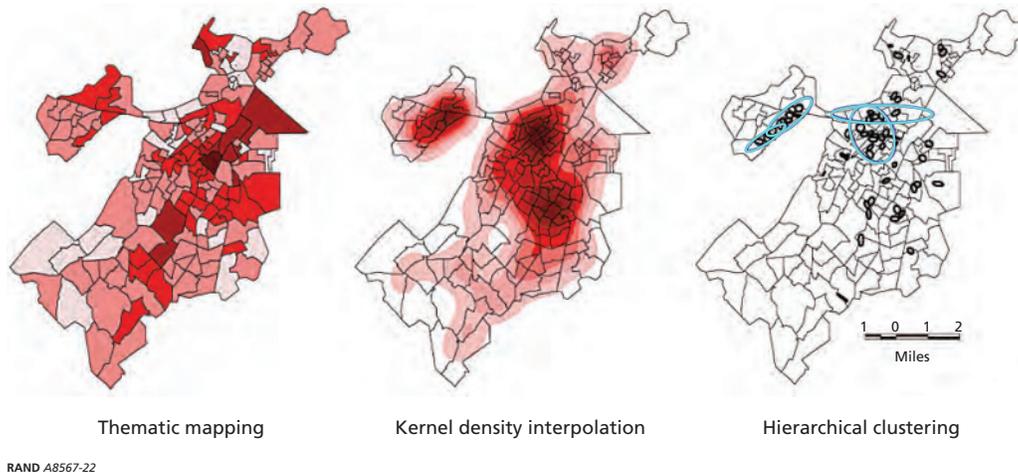
The three categories of geospatial hot spot identification techniques used here are illustrated in Figure 2.1, reflecting maps of Boston burglary events in 1999 and provided by Cameron and Leitner (2005). The first map reflects burglary rates per 100,000 residents by Census tract, the second map reflects the density per square mile, and the final is a clustering of events contained within ellipses.² Although the

¹This is not an exhaustive list of hot spot analysis categories, but it does include those approaches that use a sample of automated identification of hot spots rather than subjective visual interpretation.

²Since the purpose of this illustration is to simply compare the maps resulting from the various approaches, the density and rate scales are not shown.

unit of analysis differs among approaches, they all yield clusters of events that may be considered to be hot spots as long as the spatial dependence and concentration properties are met (see the hot spot definition provided on page 2 in Chapter One). It should be noted that the use of hierarchical clustering has been extensively discussed in spatial analysis as it is one source of the well-known *modifiable areal unit problem* (MAUP) (Openshaw, 1984) that may lead to misinterpretation of results due to the arbitrary boundaries that are used to aggregate data. Application of AHS does not resolve the MAUP, so those who interpret the results should consider that the problem may still exist.

Figure 2.1
Boston Burglary Rates, 1999



Source: Cameron and Leitner, 2005.

This chapter addresses the first research question: *Can existing geospatial tools be modified to ensure that any identified hot spots are actionable given known spatial resource constraints?* For each of the listed categories of hot spot analysis, it is possible to modify the underlying algorithms to generate clusters to allow spatial, temporal, and quantity resource constraints to be applied. We found that one particular type of hierarchical clustering method that is widely available to researchers in quantitative fields — the *complete-link method* — can be leveraged to enforce spatial constraints on

cluster sizes. We investigated various other clustering methods and determined that the complete-link method was the one best suited for enforcing spatial constraints — although other methods could be modified to yield a similar result.

Modifications to existing hot spot approaches can also be made that will allow temporal constraints to be recognized during the process of hot spot identification; by measuring the degree to which expected deployment of resources and historical problem events are synchronized in a hot spot, it is possible to provide a prioritized list of spatially constrained hot spots that may yield a more effective and efficient use of scarce resources. Finally, where quantity constraints exist, resources can be applied to the ordered list of priority hot spots until the resources are exhausted. The appropriate enhancements that need to be made to existing approaches to turn *hot spots* into *actionable hot spots* will be detailed in the remainder of this chapter. In this chapter, the example data set shown in Table 2.1 will be used to illustrate the underlying processes involved in standard hot spot identification approaches and to demonstrate how the enhancements can be applied to yield spatially actionable hot spots. This example set was carefully manufactured to be used throughout this report and to highlight how the *actionable hot spot* identification and prioritization process differs from standard approaches.

Table 2.1
Example Data, $n = 7$

Observation	Longitude	Latitude	Category	Poverty Rate	Date	Time
1	43.209	34.195	1	4.1%	12-Jun-09	07:38 PM
2	43.213	34.200	2	12.5%	09-Jun-09	11:39 PM
3	43.212	34.202	2	7.8%	05-Jun-09	08:15 AM
4	43.200	34.200	1	22.3%	21-Jun-09	10:12 PM
5	43.205	34.200	1	22.3%	23-Jun-09	02:31 AM
6	43.203	34.202	2	22.3%	18-Jun-09	01:32 PM
7	43.210	34.210	2	0.4%	01-Jul-09	09:36 AM

Each of the $n = 7$ data observations in the example contains six variables that describe the notional disorder event; two spatial variables (longitude, latitude),³ a

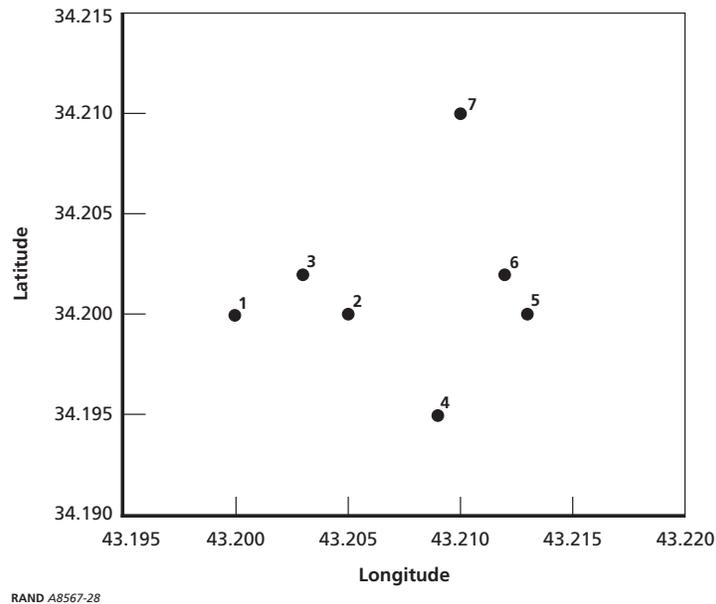
³Data may be coded using other conventions, such as the Military Grid Reference System (MGRS) or street addresses. In order to apply spatial constraints, these data must first be *geo-coded* so they are represented by latitude and longitude. In this simplified analysis, it is assumed that all observations lie on the earth — so elevation is 0 and can be ignored. While the decisions about appropriate programming environments are left to the analyst, it would be wise to select an environment that copes well with geospatial data. Using an environment that can easily interpret

variable summarizing the category of event (e.g., a type of crime, a type of improvised explosive device), one environmental variable (poverty rate), and two temporal variables (date and time the event occurred). The spatial variables can be used to plot the example data on a map (see Figure 2.2).

Hierarchical Clustering

Before introducing the spatial, temporal, and quantity constraints that make the hot spots actionable, it is first useful to provide a basic understanding of how clusters are built and hot spots are identified. Since the solution to enforcing spatial constraints lies in the employment of the complete-link method, we first focus attention on the basics of hierarchical clustering (which includes the complete-link method).

Figure 2.2
Unconstrained Clustering of Sample Data (single-link method)



latitude and longitude coordinates, compute geographic distances, and easily manipulate data arrays will save time and help to increase computational speed.

The objective of clustering is to group observations by assigning an index to each observation in a data set; the index indicates the cluster to which the observation has been assigned. Existing statistical clustering algorithms use some definition of nearness, association, or similarity⁴ between *every* pair of observations during the process of assigning a cluster index. For example, if each of seven observations in a data set — $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{6\}, \{7\}$ — is to be assigned to a cluster index, then one possible clustering outcome implied by the set of similarity values between the pairs of the seven observations may be:

- two observations are given a group index of 1 and assigned to Cluster 1 = $(\{1\}, \{2\})$
- four of the remaining observations are indexed with a 2, indicating that they belong to Cluster 2 = $(\{2\}, \{4\}, \{5\}, \{7\})$
- the remaining observation⁵ is indexed with a 3, indicating that it belongs to Cluster 3 = $(\{6\})$.

Similarity between pairs of observations is representable by a single value based on a comparison of each of the variables in an observation. It is convenient to transform measures of similarity into ones of *dissimilarity*. The transformation may be computed in any number of ways (e.g., a linear or more complicated transformation of similarity), as long as the two have an inverse relationship.

Dissimilarity

A necessary component of a dissimilarity measure is the distance function $d(\mathbf{x}_i, \mathbf{x}_j)$. A class of distance functions useful for quantitative variables in k -dimensional space, \mathbb{R}^k , is given by the Minkowski *metric distance function*:⁶

⁴For ease of discussion, the general term *similarity* will be used interchangeably with association or nearness.

⁵A cluster from a data set with n observations can be made up of anywhere between 1 and n observations.

⁶A metric distance function, $d(\mathbf{x}_i, \mathbf{x}_j)$, on all $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k \in \Omega$ returns a non-negative real value and satisfies the following properties:

$$d(\mathbf{x}_i, \mathbf{x}_j) = \left(\sum_{m=1}^k \alpha_m^\gamma |x_{i,m} - x_{j,m}|^\gamma \right)^{1/\gamma} \quad (\gamma \geq 1, \alpha_m \geq 0).^7 \quad (2.1)$$

Creating the Distance Matrix

The pairwise distances for all pairs of n vertices, $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$, can be represented in an $n \times n$ distance matrix containing a maximum of $\binom{n}{2}$ non-zero elements⁸ and can be expressed as follows:

$$\mathbf{D} = \begin{bmatrix} 0 & d(\mathbf{x}_1, \mathbf{x}_2) & \cdots & d(\mathbf{x}_1, \mathbf{x}_{n-1}) & d(\mathbf{x}_1, \mathbf{x}_n) \\ d(\mathbf{x}_2, \mathbf{x}_1) & 0 & \cdots & d(\mathbf{x}_2, \mathbf{x}_{n-1}) & d(\mathbf{x}_2, \mathbf{x}_n) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ d(\mathbf{x}_{n-1}, \mathbf{x}_1) & d(\mathbf{x}_{n-1}, \mathbf{x}_2) & \cdots & 0 & d(\mathbf{x}_{n-1}, \mathbf{x}_n) \\ d(\mathbf{x}_n, \mathbf{x}_1) & d(\mathbf{x}_n, \mathbf{x}_2) & \cdots & d(\mathbf{x}_n, \mathbf{x}_{n-1}) & 0 \end{bmatrix}.$$

Since the constraints will be imposed on the spatial variables of the data, it is convenient to first partition the data matrix, \mathbf{X} , as follows:

$$\mathbf{X} = \mathbf{S} \cup \mathbf{Y},$$

where \mathbf{S} represents the spatial variables (latitude (degrees), longitude (degrees)) and \mathbf{Y} includes all of the remaining $k' < k$ variables: $\mathbf{S} \cap \mathbf{Y} = \emptyset$.

$$\forall \mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k \in \Omega,$$

- i) **(Identity)** $d(\mathbf{x}_i, \mathbf{x}_i) = 0$
- ii) **(Positivity)** $d(\mathbf{x}_i, \mathbf{x}_j) > 0$ unless $\mathbf{x}_i = \mathbf{x}_j$ in which case $d(\mathbf{x}_i, \mathbf{x}_j) = 0$
- iii) **(Symmetry)** $d(\mathbf{x}_i, \mathbf{x}_j) = d(\mathbf{x}_j, \mathbf{x}_i)$
- iv) **(Triangle Inequality)** $d(\mathbf{x}_i, \mathbf{x}_j) \leq d(\mathbf{x}_i, \mathbf{x}_k) + d(\mathbf{x}_k, \mathbf{x}_j)$.

The space Ω , together with the metric d , is called a *metric space*.

⁷ γ is a non-negative exponent and α_m is a non-negative value indicating how much weight the individual feature m should contribute to the overall distance value. For $\gamma=1$ and $\alpha_m = 1 \forall m$, the metric is the *Manhattan distance* and for $\gamma=2$ and $\alpha_m = 1 \forall m$, the metric is the well-known *Euclidean distance*.

⁸Since this is a valid metric, the diagonal elements represented by $d(\mathbf{x}_i, \mathbf{x}_i)$ are zero (some programs have working precision limitations, so it may be necessary to force this property).

The longitude and latitude for observation i are present in the spatial data matrix, $\mathbf{S} = \{\text{longitude}_i, \text{latitude}_i\}$, and the spatial distance between pairs of points is computed using the Great Circle Distance (GCD)⁹ and represented by $d(\mathbf{s}_i, \mathbf{s}_j)$. For example, the GCD between observations 1 and 2 in the example data is represented as $d(\mathbf{s}_1, \mathbf{s}_2) = d_{1,2} = 0.4604$ km, reflecting how far apart these two observations lie on the Earth. The GCD between all pairs of observations can also be computed as $d_{1,3} = 0.2889$ km, $d_{5,6} = 0.2409$ km, etc. The full distance matrix, \mathbf{D} , for the example data representing only the spatial variables¹⁰ is

$$\mathbf{D} = \begin{bmatrix} 0 & 0.4604 & 0.3548 & 0.9983 & 1.1970 & 1.1271 & 1.4446 \\ 0.4604 & 0 & 0.2889 & 0.6674 & 0.7366 & 0.6819 & 1.2047 \\ 0.3548 & 0.2889 & 0 & 0.9552 & 0.9473 & 0.8286 & 1.0993 \\ 0.9983 & 0.6674 & 0.9552 & 0 & 0.6674 & 0.8268 & 1.6724 \\ 1.1970 & 0.7366 & 0.9473 & 0.6674 & 0 & 0.2409 & 1.1470 \\ 1.1271 & 0.6819 & 0.8286 & 0.8268 & 0.2409 & 0 & 0.9094 \\ 1.4446 & 1.2047 & 1.0993 & 1.6724 & 1.1470 & 0.9094 & 0 \end{bmatrix}.$$

The GCD is used for three reasons:

1. Computing distances between the latitude and longitude variables expressed in degrees leads to distortions, since the spacing between degrees latitude depends on where the observation lies on the Earth (a one-degree separation is much smaller near the poles than near the Equator).
2. The spatial constraints of the resources to be applied later are normally measured in miles or kilometers, not degrees, so a common distance scale is required.

⁹Let R = Earth's radius (mean radius = 6,371 km)

$\Delta\text{lat} = (\text{lat}_i - \text{lat}_j)$ (the difference in degrees latitude)

$\Delta\text{lon} = (\text{lon}_i - \text{lon}_j)$ (the difference in degrees longitude)

Then $GCD = R \cdot \arctan \left(\frac{\sqrt{[\cos(\Delta\text{lat}) \sin(\Delta\text{lon})]^2 + [\cos(\text{lat}_i) \sin(\text{lat}_j) - \sin(\text{lat}_i) \cos(\text{lat}_j) \cos(\Delta\text{lon})]^2}}{\sin(\text{lat}_i) \sin(\text{lat}_j) + \cos(\text{lat}_i) \cos(\text{lat}_j) \cos(\Delta\text{lon})} \right)$.

¹⁰Clustering may be executed using any weighted subset of $m = 1, \dots, k'$ variables, but the geographic constraint need only be applied to the spatial variables. A revised version of the Minkowski distance metric for any pair of points in X may then be represented by

$$d(\mathbf{x}_i, \mathbf{x}_j) = \left(\alpha_o d(\mathbf{s}_i, \mathbf{s}_j)^\gamma + \sum_{m=1}^{k'} \alpha_m |y_{i,m} - y_{j,m}|^\gamma \right)^{1/\gamma} \quad (\gamma \geq 1, \{\alpha_o, \alpha_m\} \geq 0). \quad (2.2)$$

3. It is a valid metric with properties that ease the computational burden associated with clustering.

Hierarchical Clustering with Spatial Constraints

Among hierarchical clustering algorithms,¹¹ there exist *divisive* and *agglomerative* approaches. Divisive algorithms begin with all observations belonging to a single cluster and then, guided by an explicit division rule, iteratively partition the observations into smaller clusters, with each observation normally belonging to one-and-only-one cluster.¹² Hierarchical agglomerative clustering methods (HACMs) begin with an initial set of n clusters — each containing one of the n observations in the data set — and then use a decision rule to iteratively merge clusters until the single, final cluster contains all n observations. In this report, agglomerative clustering is used, since its application to the resource-constraint problem is both more intuitive and easier to implement. While there are several types of agglomerative hierarchical clustering decision rules (also called “methods”) — single-link method, complete-link method, Ward’s method, K -means method, centroid method, nearest neighbors method, etc. (Gordon, 1999) — only the complete-link method allows strict enforcement of spatial constraints on the cluster geometry that reflect the physical limits of the resource, although other methods may be modified to yield the same solution.

The complete-link and single-link methods represent the two extremes of decision rules that may be applied in the process of agglomerative clustering. Other methods may yield clusters that also differ from the complete-link method, but the single-link method has been chosen as an illustration. For each pair of distinct clusters in an environment where the range of the resource has radius r (measured in kilometers), and $\{C_p, C_q\}$ are two clusters, the decision rules specifying whether merging of clusters can occur are

$$\text{Single – linkmethod : } \min\{d(C_p, C_q)\} \leq 2r \quad (2.3)$$

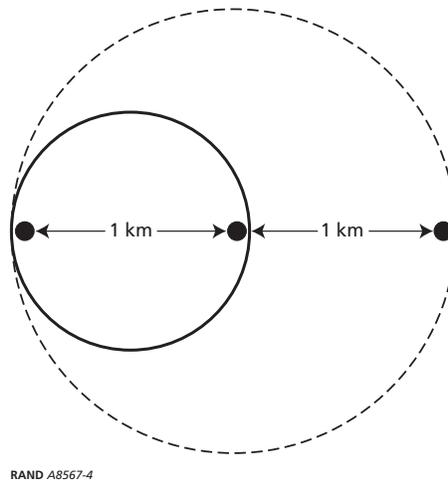
$$\text{Complete – linkmethod : } \max\{d(C_p, C_q)\} \leq 2r. \quad (2.4)$$

¹¹Aside from hierarchical clustering algorithms, other families of algorithms exist (Gordon, 1999), but since none of these families possesses the properties that will allow spatial resource constraints to be applied, they will not be discussed.

¹²There are algorithms that allow observation-sharing by clusters. In Chapter Three, we detail an algorithm that allows observation-sharing.

In the single-link method, this means that two clusters may be merged if the distance between *any* observation in \mathcal{C}_p and *any* observation in \mathcal{C}_q is less than the diameter ($s = 2r$) of the maximum range (represented as a circle) of the resource that may be deployed to the cluster. To see how this method is inappropriate for identifying clusters against which resources may be deployed, suppose there are three vertices along a straight line separated by 1 km (see Figure 2.3). If the resource has a radius of $r = 0.5$ km, the single-link method would allow all three vertices to be merged into one cluster, since the minimum distance between any pair of vertices is within the allowable diameter ($s = 2r = 1$ km). The dotted-line circle represents the size of the cluster allowed by the single-link method and the solid-line circle represents the maximum size of the resource. Therefore, unless explicitly modified to test to enforce the spatial constraint, the single-link clustering method allows clusters that exceed the range of the constrained resource and are not fully actionable.

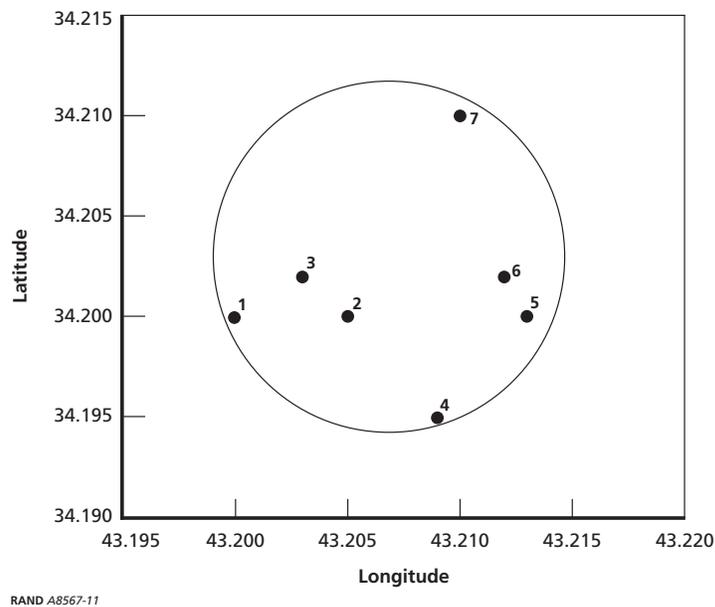
Figure 2.3
Example of Single-Link Clustering



On the other hand, the complete-link algorithm would not allow all three observations to be contained in a single cluster and would have a diameter no greater than the limited range of the resource. This enforcement of the limits of the resource guarantees that all clusters identified are spatially actionable given a specified constraint. This is because the rule for merging clusters can directly reflect the footprint of the

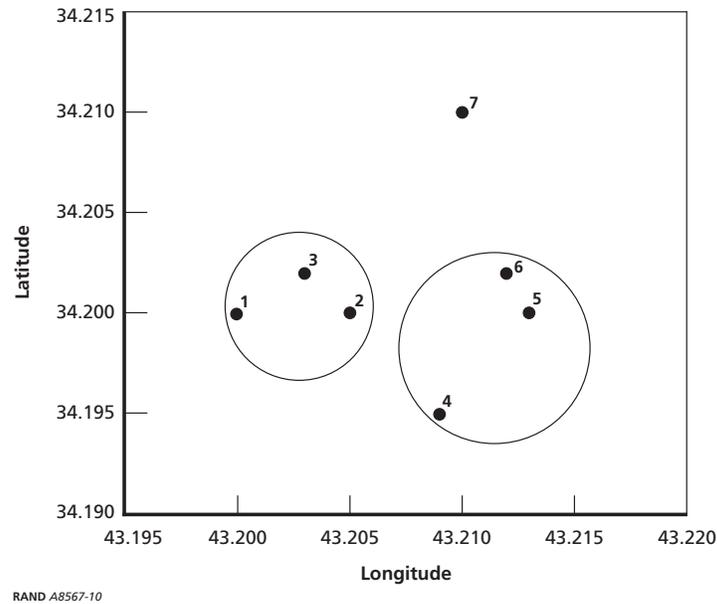
resource. The rule is that two clusters may be merged if-and-only-if the distance between *every* observation in one cluster and *every* observation in the cluster with which it is to be merged is no greater than $s = 2r$. This ensures that every observation within the cluster falls within the footprint of the constrained resource (if the footprint is circular); the results from application of the complete-link method would appear as the solid circle in Figure 2.3. The final clusters that would result for the example data using the single-link and complete-link methods are shown in Figures 2.4 and 2.5, respectively.

Figure 2.4
Clustering Results (single-link method)



Other HACM algorithms that employ cluster-merging decision rules, such as Ward's method, nearest neighbors, and the K -means method (Gordon, 1999), do not easily allow a strict limit on the diameter of the resource to be enforced. It is possible that these methods may yield the same clustering outcomes for certain data sets that respect known spatial resource constraints, but only the complete-link method guarantees this condition will be met for all data sets. It may also be possible to accomplish this result using nonhierarchical clustering techniques, but it was our

Figure 2.5
Constrained Clustering Results (complete-link method)



aim to suggest how spatial constraints may be applied when one or more of the three commonly used categories of hot spot identification tools (thematic mapping, kernel density interpolation, and hierarchical clustering) are selected by the user. For that reason, we make the following recommendation: **Enforcement of spatial constraints on clusters that are eligible to be identified as “hot spots” can be achieved with certainty if the *complete-link method* of hierarchical agglomerative clustering is used; without modification, other hierarchical clustering methods may not guarantee this outcome.**

Spatial Dependence

Establishing spatial dependence is one requirement for determining that a group of observations is indeed a cluster (and also for labeling as “hot spots,” since these are special cases of clusters), but the hierarchical clustering method described above makes no explicit mention of this type of test. Before considering that an observation may join another cluster, it must be determined that it is closer than would be

expected if the spacing had occurred randomly. Existing approaches apply the Nearest Neighbor Index (NNI) test to accomplish this. The NNI compares the spatial distances among the n observations in the study area with n randomly spaced observations in an area of the same size. The distance between an observation and other observations in the data set divided by the average distance of the observations in the randomly generated data set yields the NNI.

Values less than 1.0 indicate spatial dependence and allow the observation to be considered for joining clusters — although a formal hypothesis may require that the values be below a threshold smaller than 1.0. To ensure consistency with the complete-link method, the comparison needs only a slight modification: The comparison of actual observations with others in the data set should only include distances that are no greater than $s = 2r$. This simple modification essentially ensures that identification of observations that may join clusters is subject to the same spatial constraints that will be enforced when clusters are actually built. For that reason, we make the following recommendation: **When spatial constraints are to be applied in the building of clusters, tests for spatial dependence should also be modified to reflect the constraints.**

Augmenting Other Hot Spot Identification Approaches

Since the complete-link method is a form of hierarchical clustering used in hot spot identification, the modifications required to enforce spatial constraints — and thereby to ensure actionability — are minor. The decision rules for merging clusters need only be updated to ensure that the resulting clusters do not exceed the range of the resource. If direct modification cannot be made to existing proprietary algorithms, this enforcement of spatial constraints can still be accomplished by taking the distance matrix that emerges from the user's software of choice and passing it through any statistical software (e.g., R or SAS) that is capable of applying the complete-link method. Another type of clustering available for building hot spots in some geospatial software packages is the K -means clustering method. This approach (Chainey et al., 2002) partitions the data into a user-defined number (K) of groups and encloses them with ellipses. Note that this approach, although useful in some types of analysis, is inconsistent with problems where resources are constrained, since the size of the ellipses is unbounded by design.

Spatial constraints can also be applied when the kernel density approach is the

preferred method for building clusters.¹³ However, since that approach essentially spreads the observations over wider areas, the modifications required to implement spatial constraints are more elaborate. Rather than representing observations as points on a map, they are represented by shaded grid squares that surround the location of the actual event. Figure 2.6 illustrates one example of how a single observation is spread over several adjacent grid squares.

To apply spatial constraints, the location of each grid square must first be extracted. Then, instead of using the actual observations to build clusters, the coordinates of the grid squares are used during the clustering process. Assuming that the location of the grid square is represented by its center, clusters are built among the pseudo-observations (the grid square centers) instead of the actual observations. This is arguably a more computationally intensive exercise, yet it is required to ensure that the spatial constraints are enforced with the complete-link method. Essentially, this means that the building of clusters requires two separate analyses of the data: one to apply the KDE and another to cluster the resulting grid squares using the complete-link HACM.¹⁴

The final category of hot spot methods that needs to be addressed is thematic

¹³*Kernel density estimation* (KDE) is an increasingly popular method for visualizing spatial data and identifying hot spots. This interpolation technique creates a relatively smooth, continuous, color-coded surface that represents the number of events across the area. The basic mechanics of this approach require that, rather than representing a historical event as a single observation on a map, the event is spread evenly over a predefined area surrounding the actual event. The aim is to alleviate the difficulty associated with visually interpreting areas with higher concentrations of disorder events across the study area. The size of the surrounding area and the way that the density of the actual spatial observation is allocated over that area are determined by a mathematical function called a *kernel*:

$$k(\mathbf{s}; b), b \geq 0 \text{ such that } \int k(\mathbf{s}; b) d\mathbf{s} = 1.$$

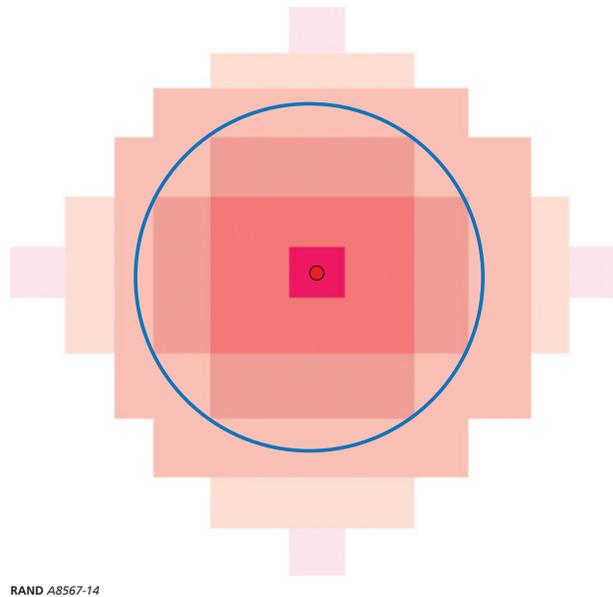
One familiar example is the normal kernel:

$$k(\mathbf{s}; b) = \frac{1}{\sqrt{2\pi}} e^{-\frac{\mathbf{s}^2}{2b}},$$

which is a valid probability distribution that allocates the entirety of the event's mass over a larger area. The value of b is known as the *bandwidth* and indicates the size of the area over which the density should be allocated.

¹⁴One additional modification is needed to ensure that the entire grid square is entirely enclosed within the circle during the application of the HACM pass over the data. The effective radius will need to be reduced by $g\sqrt{2}$, where g is the half-width of the grid square. This adjustment corrects for the coarseness that results from gridding the data.

Figure 2.6
KDE Representation of a Spatial Observation



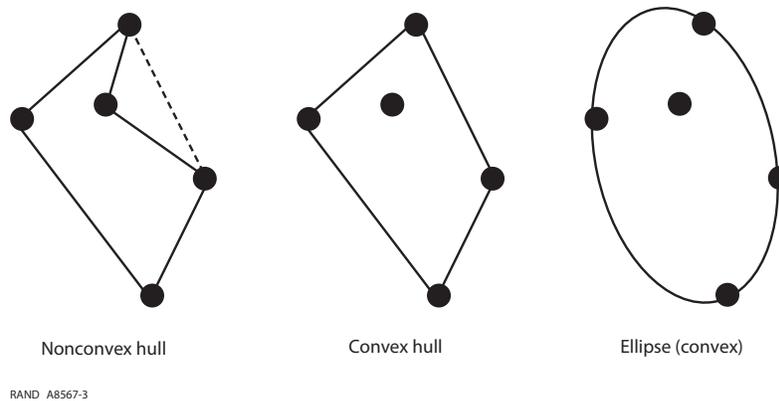
mapping. Clearly, unless the administrative boundaries of an area (or clusters of adjacent areas) fit inside a circle of radius r , the spatial constraint will not be recognized. For that reason, identification of actionable hot spots is possible for users of the thematic mapping approach in only a limited number of cases. When the partitions (or clusters of adjacent partitions) do fit neatly within the footprint of the deployable resource, application of our approach would require that a catchment polygon (an object surrounding the entire administrative boundary) first be created using either an ellipse or convex hull.¹⁵ Figure 2.7 illustrates three types of catchment polygons (from left to right): The first is a nonconvex hull,¹⁶ the middle diagram is a valid convex hull, and the last diagram is an ellipse (also convex).

In summary, existing approaches to hot spot identification can be modified to allow

¹⁵A *convex hull* is defined as a polygon surrounding data points in which any line that can be drawn between observations in the hull does not fall outside the polygon.

¹⁶It is nonconvex since a line between one pair of points lies outside the polygon.

Figure 2.7
Catchment Methods for Tests of Clustering



spatial constraints to be enforced so that the resulting areas under consideration for application of resources by decisionmakers are appropriately sized to deter, disrupt, or prevent future disorder events.

Spatially Constrained Resources with Noncircular Footprints

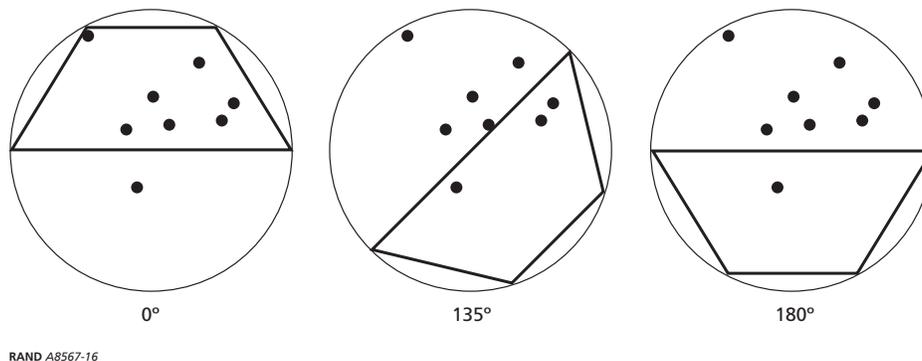
In the description of the AHS approach discussed thus far in this report, there has been an explicit assumption that the resources have a circular footprint. Since that is rarely the case, we now discuss application of spatial constraints when resource footprints are noncircular. The footprint of the resource may be represented as a polygon (convexity is not required; for example, some airborne assets project an “axe-blade”-shape footprint on the Earth). A simple example of a polygon footprint would be a trapezoid that may result from a stationary, elevated surveillance camera. Before applying the complete-link HACM (or modified versions of the KDE or thematic mapping approaches), the effective radius must be found. This is done by computing the maximum distance between all pairs of points that define the polygon footprint. This maximum value, $2r_{max}$, becomes the effective diameter of the resource, which is

used in the application of the complete-link HACM.

Once the clusters have been built using the complete-link HACM (or modified versions of the KDE or thematic mapping approaches) assuming a circular footprint with effective diameter $2r_{max}$, a resource with a noncircular footprint may be superimposed over the cluster to check that all of the observations fall within the polygon footprint. Figure 2.8 demonstrates how the selection of observations to be included in the cluster would be identified. The center of the minimum volume enclosing circle and the score-weighted cluster centroid are chosen as candidate rotation points. For each rotation point, the following steps are executed:

1. The polygon is centered on the rotation point at an arbitrary angle.
2. A sub-cluster is identified that contains all observations that lie inside the superimposed polygon.
3. The polygon is rotated by a small, fixed amount about the rotation point and another sub-cluster is identified.
4. The rotation continues until the polygon is in its original position (Figure 2.8 demonstrates the effect of three different rotations [0, 135, and 180 degrees]).
5. Each sub-cluster containing at least two observations is identified as a candidate hot spot.

Figure 2.8
Polygon Spatial Constraint Fitting



This approach is most useful when the analyst is looking for clusters over neighborhoods or large areas. When the level of analysis involves a linear search area (such as identifying linear streets with a large concentration of disorder events), this approach is inefficient. A simpler approach for linear-type searches would be to first build linear clusters using the algorithms supplied in geospatial software and superimposing rectangles over the resulting linear clusters. Finally, since it is computationally costly to search for hot spots that fit within noncircular footprints using the rotation method we described, extensions of our research will investigate more-efficient methods for addressing this problem.

Data Reduction

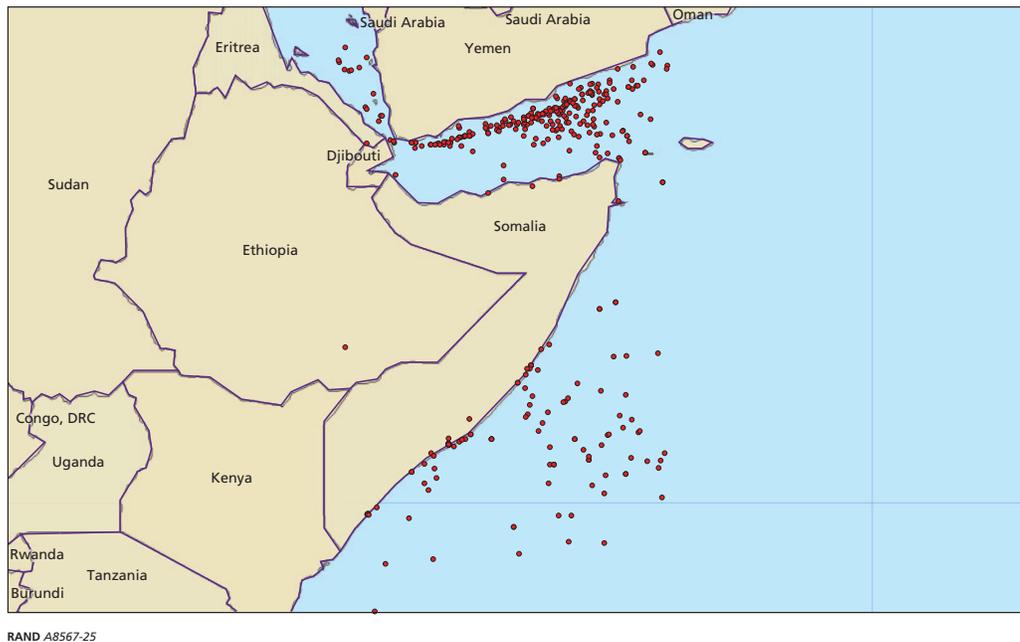
The extensions to existing approaches described so far require a significant increase in computational burden to execute. One solution — closely linked to the application of temporal constraints to be described in the next chapter — is to remove observations that are irrelevant to the analysis. A leaner data set results in far fewer computations (calculation of pairwise distances, assessing merges, etc.) and greater efficiency. Observations may be removed from the analysis for three reasons:

Temporal. In some research areas, there may be some reason to believe that some historical observations are simply too old to be included in the analysis. For example, Figure 2.9 shows the historical piracy events in the Gulf of Aden (GoA) and coastal Somalia between 2004 and 2008, while Figure 2.10 shows the piracy events in the same region for only July through December 2008. Clearly, the pirates have begun to demonstrate a more recent preference for conducting attacks off the coast of Yemen. For that reason, identification of actionable piracy hot spots should take this shift in preferences into account and reduce the weight of — or omit — observations that are more indicative of older spatial patterns in piracy attacks.

Categorical. Some observations in the data set may pertain to phenomena that are not being considered for resource allocation by the decisionmaker with scarce resources. For example, if the objective of the deployment of resources is to reduce cases of burglary, perhaps the observations that indicate cases of homicide can be omitted from the analysis (unless a correlation between crime categories in certain areas is a useful piece of information).

Spatial. With a spatial constraint implied by an actionable resource radius of r and the knowledge that a cluster must have more than one observation to be considered a hot spot, an observation that is not within $s = 2r$ of any other observation *cannot be part of a hot spot* and can be removed. Since this requires that the pairwise distances be computed first, the gains in computational efficiency are less than in cases where data are removed for other reasons.

Figure 2.9
Piracy Incidents in GoA/Somali Coast, 2004 – 2008

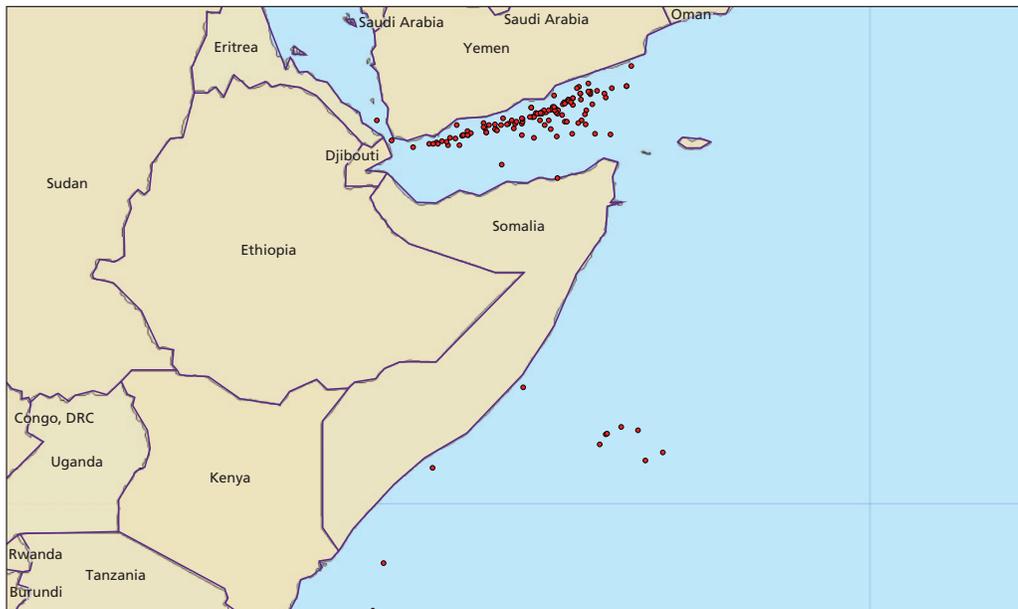


Source: National Geospatial Intelligence Agency.

Summary

This chapter has demonstrated that the combination of existing statistical methods and our new innovations can be used to identify hot spots that are spatially actionable. Hot spot identification methods used by the geospatial analysts include a broader set of approaches that are selected based on their appropriateness for understanding the

Figure 2.10
Piracy Incidents in GoA/Somali Coast, July – December 2008



RAND A8567-26

Source: National Geospatial Intelligence Agency.

underlying problems. However, it has been shown that our approach can augment these existing approaches and allow the user to both employ his/her tool of choice *and* enforce spatial constraints. So, with enhancements, the geospatial analyst can identify hot spots based on the needs of the end-user and help guide resource allocation decisions.

Cluster Point-Sharing

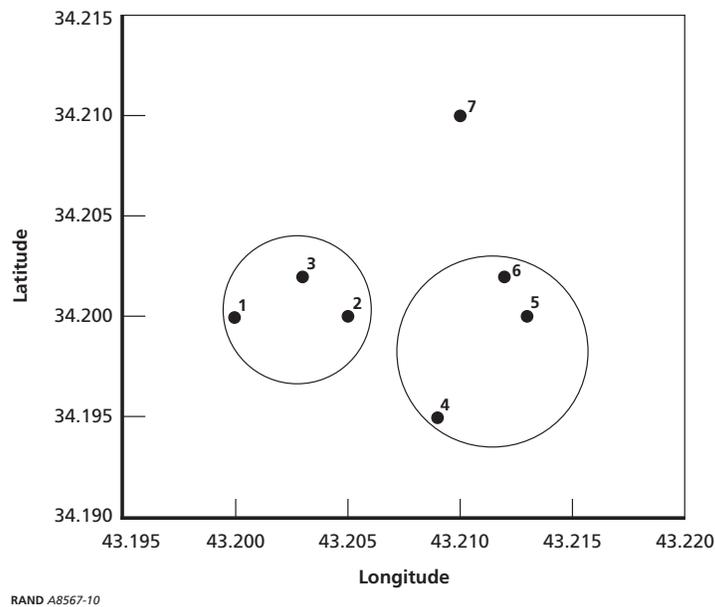
Due to the computational burden associated with hierarchical clustering, most existing algorithms yield suboptimal solutions. For data sets with n data points or *observations*, the operations that would be required to search over all possible clustering outcomes for an optimal solution is proportional to n -factorial ($n!$). Even with relatively small n , the required computing time becomes unmanageable and shortcuts need to be taken. For that reason, known implementations of hierarchical clustering use a “one-step ahead” strategy that results in the inability to reverse merges. These implementations iteratively build clusters with the stipulation that, once an observation joins a cluster, it cannot be separated from that cluster throughout future iterations. The merges of clusters are chosen to be the best during the iteration in which they occur, but the inability to reverse a decision limits the algorithms’ ability to provide a “better” clustering as more information about the structure of the data emerges in later iterations.

To mitigate the problems associated with the one-step-ahead approach, an algorithm was created to perform a *second pass* through the clusters generated by the complete-link HACM, to determine whether additional observations could be extracted from other clusters. With a resource that is constrained to have a diameter of $s = 2r$, the basic idea is to expand clusters to include observations belonging to other clusters as long as the constraint on the diameter is not violated. This will allow the analyst to build clusters that contain more observations and, hence, are considered “hotter” spots. The end result is that the identified clusters represent a better summary of the underlying patterns in the data rather than the artificiality of the algorithms used to generate them. The approach allows observations to be shared by more than one cluster. To some observers, this may violate the notion (and hence the purpose) of clustering. However, for some applications, it is likely

that these observations may indeed have relevance to the analysis of both clusters to which they may belong. If the intent is to build actionable clusters (particularly those defined as “hot spots”), the proposed approach provides a reasonable balance between computational complexity and utility.

To illustrate how the algorithm functions, we again use the same example data set that we used in Chapter Two. The complete-link HACM yielded clusters (shown in Figure 3.1) that respect the physical constraints of the resource with maximum radius r . The algorithm begins with this result and attempts to find a clustering solution that shares observations among clusters while preserving the spatial constraint.

Figure 3.1
Constrained Clustering Results (complete-link method)



The Second-Pass Approach

This section describes the mechanics of the two-pass algorithm. The first pass applies the HACM and yields a set of clusters in which each observation belongs to only one of the identified clusters.

Suppose that the original data set contained n spatial observations and that, in the course of the complete-link clustering process — and subsequent *cutting* to consider only those clusters with a maximum internal distance of $s = 2r - m$ observations ($m \leq n$) joined a cluster. The first step in the process is to sort the data set, \mathbf{X} , so that the first m rows of the Θ_c matrix reflect the m observations in cluster c . Based on the results from the first-pass clustering of the example data shown, the resulting cluster assignments are $\mathcal{C}_8 = \{\mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6\}$, $\mathcal{C}_9 = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$, and $\mathcal{C}_7 = \{\mathbf{x}_7\}$. Focusing on \mathcal{C}_8 , the sorted data matrix, \mathbf{X}'_8 , is shown in Table 3.1.

Table 3.1
Example Data Resorted for Cluster 8

Cluster	Sorted Obs.	Original Obs.	Longitude	Latitude	Category	Poverty Rate	Date	Time
8	4	1	43.209	34.195	1	4.1%	12-Jun-09	07:38 PM
8	5	2	43.213	34.200	2	12.5%	09-Jun-09	11:39 PM
8	6	3	43.212	34.202	2	7.8%	05-Jun-09	08:15 AM
9	1	4	43.200	34.200	1	22.3%	21-Jun-09	10:12 PM
9	2	5	43.205	34.200	1	22.3%	23-Jun-09	02:31 AM
9	3	6	43.203	34.202	2	22.3%	18-Jun-09	01:32 PM
7	7	7	43.210	34.210	2	0.4%	01-Jul-09	09:36 AM

For the c^{th} cluster, let Θ_c describe the relationship between each of the m observations in cluster c and each of the observations in all other clusters. Then, for all i, j in n , let $\theta(\mathbf{x}'_i, \mathbf{x}'_j)$ be a binary operator applied to the sorted data matrix¹ \mathbf{X}' , where

$$\theta(\mathbf{x}'_i, \mathbf{x}'_j) = \begin{cases} 0 & \text{if } d(\mathbf{x}'_i, \mathbf{x}'_j) > 2r \\ 1 & \text{if } d(\mathbf{x}'_i, \mathbf{x}'_j) \leq 2r. \end{cases}$$

The resulting matrix of values after application of the θ operator is:

$$\Theta_c = \left[\begin{array}{ccc|ccc} \theta(\mathbf{x}'_1, \mathbf{x}'_1) & \cdots & \theta(\mathbf{x}'_1, \mathbf{x}'_m) & \theta(\mathbf{x}'_1, \mathbf{x}'_{m+1}) & \cdots & \theta(\mathbf{x}'_1, \mathbf{x}'_n) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \theta(\mathbf{x}'_m, \mathbf{x}'_1) & \cdots & \theta(\mathbf{x}'_m, \mathbf{x}'_m) & \theta(\mathbf{x}'_m, \mathbf{x}'_{m+1}) & \cdots & \theta(\mathbf{x}'_m, \mathbf{x}'_n) \\ \hline \theta(\mathbf{x}'_{m+1}, \mathbf{x}'_1) & \cdots & \theta(\mathbf{x}'_{m+1}, \mathbf{x}'_m) & \theta(\mathbf{x}'_{m+1}, \mathbf{x}'_{m+1}) & \cdots & \theta(\mathbf{x}'_{m+1}, \mathbf{x}'_n) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \theta(\mathbf{x}'_n, \mathbf{x}'_1) & \cdots & \theta(\mathbf{x}'_n, \mathbf{x}'_m) & \theta(\mathbf{x}'_n, \mathbf{x}'_{m+1}) & \cdots & \theta(\mathbf{x}'_n, \mathbf{x}'_n) \end{array} \right].$$

¹Since the data have been sorted, the index reflects the observation number of the sorted data. For example, observation \mathbf{x}_4 is now \mathbf{x}'_1 .

Θ_c may be decomposed into four sub-matrixes that will be discussed separately below.

$$\left(\begin{array}{c|c} \mathbf{A}(c) & \mathbf{B}(c) \\ \hline \mathbf{B}^T(c) & \mathbf{E}(c) \end{array} \right)$$

Sub-matrix $\mathbf{A}(c)$

A is an $m \times m$ sub-matrix that summarizes the distances between observations in the cluster on which the focus is applied — the c^{th} cluster. Since the first pass of clustering necessarily includes only those observations that are within $s = 2r$ of all other observations in the cluster, all of the elements are equal to 1. For example, the distance matrix of Great Circle Distances for \mathbf{X}'_8 is

$$\mathbf{D}'_8 = \begin{bmatrix} 0.0000 & 0.6674 & 0.8268 & 0.9983 & 0.6674 & 0.9552 & 1.6724 \\ 0.6674 & 0.0000 & 0.2409 & 1.1970 & 0.7366 & 0.9473 & 1.1470 \\ 0.8268 & 0.2409 & 0.0000 & 1.1271 & 0.6819 & 0.8286 & 0.9094 \\ 0.9983 & 1.1970 & 1.1271 & 0.0000 & 0.4604 & 0.3548 & 1.4446 \\ 0.6674 & 0.7366 & 0.6819 & 0.4604 & 0.0000 & 0.2889 & 1.2047 \\ 0.9552 & 0.9473 & 0.8286 & 0.3548 & 0.2889 & 0.0000 & 1.0993 \\ 1.6724 & 1.1470 & 0.9094 & 1.4446 & 1.2047 & 1.0993 & 0.0000 \end{bmatrix},$$

and so the Θ_8 matrix is given by:

$$\Theta_8 = \left(\begin{array}{ccc|ccc} 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ \hline 1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 \end{array} \right),$$

of which sub-matrix A is

$$\mathbf{A} = \begin{bmatrix} \theta(x'_1, x'_1) & \theta(x'_1, x'_2) & \theta(x'_1, x'_3) \\ \theta(x'_2, x'_1) & \theta(x'_2, x'_2) & \theta(x'_2, x'_3) \\ \theta(x'_3, x'_1) & \theta(x'_3, x'_2) & \theta(x'_3, x'_3) \end{bmatrix} = \begin{bmatrix} \theta(x_4, x_4) & \theta(x_4, x_5) & \theta(x_4, x_6) \\ \theta(x_5, x_4) & \theta(x_5, x_5) & \theta(x_5, x_6) \\ \theta(x_6, x_4) & \theta(x_6, x_5) & \theta(x_6, x_6) \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

Sub-matrixes $\mathbf{B}(c)$, $\mathbf{B}^T(c)$

Sub-matrix $\mathbf{B}(c)$ yields the same information as $\mathbf{B}^T(c)$. Since the matrix is symmetric, the elements are simply transpositions of one another, so only $\mathbf{B}(c)$ will be discussed. The binary value in the i^{th} row and j^{th} column of $\mathbf{B}(c)$ indicates whether the GCD between observation i in the c^{th} cluster and observation j are within $s = 2r$ of one another. For each of the j ($j = 1, \dots, (n - m)$) columns of $\mathbf{B}^T(c)$, a sum can be computed:

$$S_c(j) = \sum_{i=1}^{n-m} b_{i,j}.$$

If $S_c(j) = m$, the GCD between the j^{th} observation and *every* observation in \mathcal{C}_c is no greater than $s = 2r$. So, the observation may join the cluster without violating the spatial constraint. In \mathcal{C}_8 , $S_8(2)$ and $S_8(3)$ both equal m and so the cluster $\mathcal{C}_8 = \{\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6\}$ is *potentially* a cluster with a diameter no greater than $s = 2r = 1$ km and may therefore be actionable given the spatial constraint. However, an additional requirement for *both* observations \mathbf{x}_2 and \mathbf{x}_3 to join \mathcal{C}_8 is that they be within $s = 2r$ of one another — it is for this reason that they are only *potential* additions to the cluster. The final sub-matrix, $\mathbf{E}(c)$, indicates whether this additional requirement is met.

Sub-matrix $\mathbf{E}(c)$

Sub-matrix $\mathbf{E}(c)$ summarizes the GCD between every observation that has not been assigned to \mathcal{C}_c during the first pass. For each of the observations in the set identified to be a potential addition to \mathcal{C}_c , based on the operator $S_c(j)$ applied to the columns of $\mathbf{B}(c)$, it is now possible to determine which subsets can be added to \mathcal{C}_c without violating the distance constraint implied by radius r . If \mathbf{K} is the set of h observations identified as potential additions to \mathcal{C}_c , then let $\mathbf{E}'(c)$ be the sub-matrix of $\mathbf{E}(c)$ containing only those rows and columns pertaining to the elements of \mathbf{K} . For example, potential additions to \mathcal{C}_8 were identified — by the process applied to $\mathbf{B}(c)$ — to include \mathbf{x}_2 and \mathbf{x}_3 . Since

$$\mathbf{E}_8 = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

\mathbf{E}'_8 includes only those rows and columns of \mathbf{E}_8 that summarize the distances between \mathbf{x}_2 and \mathbf{x}_3 :

$$\mathbf{E}'_8 = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

There are 2^h distinct patterns of 0's and 1's generated by concatenating the rows of \mathbf{E}'_c (for example, both rows of \mathbf{E}'_8 yield pattern 11, so there is only one distinct pattern — the other possible patterns are 00, 01, and 10).

Selecting Observations for Sharing

For each distinct pattern yielded by \mathbf{E}'_c , the next step of the two-pass algorithm requires that a new cluster be built by merging \mathcal{C}_c with the other observations indicated by a value of 1 in the pattern. For example, suppose that three observations ($\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$) have been identified via the summation operator $S_c(j)$ to be potential additions to \mathcal{C}_c and that

$$\mathbf{E}'_c = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

yielding distinct patterns 110, 111, and 010. Then three possible clusters that respect the spatial limits on the range of the resource are possible:

- $\mathcal{C}_c^* = \{\mathcal{C}_c, \mathbf{x}_i, \mathbf{x}_j\}$
- $\mathcal{C}_c^{**} = \{\mathcal{C}_c, \mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k\}$
- $\mathcal{C}_c^{***} = \{\mathcal{C}_c, \mathbf{x}_j\}$.

Typically, the one with the most new observations (\mathcal{C}_c^{**} in the example) would be the most appealing. When it becomes necessary to choose between two equally sized clusters, the cluster that produces the highest score after weighting (for instance, if time is the determining factor of the weighting function, then more recent observations would be favored over older ones) is chosen (see Chapter Two for a discussion of weighting). In order to yield the two-pass cluster with the most observations, the list of distinct patterns should be sorted and processed in descending order according to the pattern with the most 1's. Going back to the example of \mathcal{C}_8 , there is only one distinct pattern, 11, indicating that both \mathbf{x}_2 and \mathbf{x}_3 are within $s = 2r$ of one another

and also within $s = 2r$ of every observation in \mathcal{C}_8 . So the second-pass algorithm yields a cluster with five observations $(\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6)$ instead of the three observations found in the first pass. Hence the second-pass algorithm allows the analyst to identify clusters with more activity that may still be actionable given the spatially constrained resource.

Iteration over All Clusters Identified in the First Pass

For every cluster identified in the first pass containing more than one observation, this process should be repeated. Clusters with one observation need not be processed since, if that observation could have joined a cluster during the first pass, it would have already been merged. Completing the example of the example data set, the second pass would be run on the cluster $\mathcal{C}_9 = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$.

The resorted data matrix is

$$\mathbf{D}'_9 = \begin{bmatrix} 0 & 0.4604 & 0.3548 & 0.9983 & 1.1970 & 1.1271 & 1.4446 \\ 0.4604 & 0 & 0.2889 & 0.6674 & 0.7366 & 0.6819 & 1.2047 \\ 0.3548 & 0.2889 & 0 & 0.9552 & 0.9473 & 0.8286 & 1.0993 \\ 0.9983 & 0.6674 & 0.9552 & 0 & 0.6674 & 0.8268 & 1.6724 \\ 1.1970 & 0.7366 & 0.9473 & 0.6674 & 0 & 0.2409 & 1.1470 \\ 1.1271 & 0.6819 & 0.8286 & 0.8268 & 0.2409 & 0 & 0.9094 \\ 1.4446 & 1.2047 & 1.0993 & 1.6724 & 1.1470 & 0.9094 & 0 \end{bmatrix},$$

and the matrix of binary indicators of eligibility for joining \mathcal{C}_9 is:

$$\Theta_9 = \left[\begin{array}{ccc|cccc} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ \hline 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{array} \right].$$

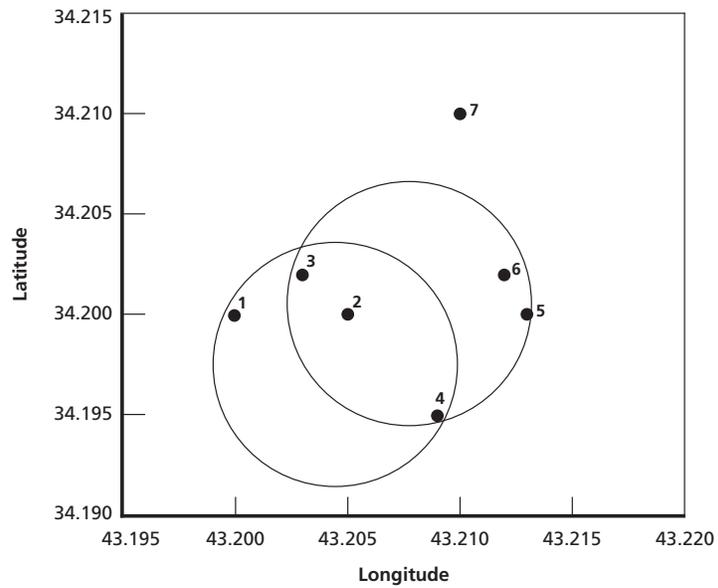
This indicates that \mathbf{x}_4 may join the cluster without violating the resource constraint yielding the final clusters:

- $\mathcal{C}_7 = \{\mathbf{x}_7\}$
- $\mathcal{C}^*_8 = \{\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6\}$

$$\bullet C^*_9 = \{x_1, x_2, x_3, x_4\}.$$

The final two-pass clustering results are shown in Figure 3.2.

Figure 3.2
Final Improved Constrained Clustering



RAND A8567-12

Hot Spot Prioritization and Performance Measurement

After imposing the spatial constraints and establishing that spatial dependence is present in identified clusters, the user will likely be left with many clusters. Three natural questions arise:

1. Which clusters are hot spots?
2. Which clusters are hotter than others?
3. Given the resource quantity constraints, against which hot spots should resources be deployed to yield maximum benefit?

This chapter answers the second research question: *Can identified actionable hot spots be prioritized so that the decisionmaker can efficiently allocate scarce resources to yield maximum effectiveness against problem areas?* In the first part of the chapter, we determine which spatially constrained clusters are actually hot spots. Following that determination, we prioritize the remaining actionable hot spots according to how well the historical events are synchronized with the expected resource deployment that is subject to temporal constraints. In the last part of this chapter, we present a calibration process that helps the analyst exploit identifiable patterns in the disorder activity to propose resource deployment solutions that are expected to yield maximum effectiveness against the problem being addressed.

From Actionable Clusters to Actionable Hot Spots

To be considered a hot spot, we earlier noted that a cluster must have two properties:

- Spatial dependence (indicating that the observations in the cluster are related to each other) must be established using statistical testing; with a reasonable amount of confidence, it can be determined that the clustering pattern could not have occurred randomly.
- The concentration of problem events in the cluster is greater than the average concentration of events in other parts of the study area.

Since the first property has already been established during the selection of observations that joined the spatially constrained clusters, for a cluster to be considered an actionable hot spot, we need only establish that the concentration of events in the clusters is greater than other parts of the study area. The standard approach to establishing a large concentration would calculate two concentration values

1. Across the study area — the total number of events in the study area is divided by the total size of the area size (in square kilometers or miles) and the resulting concentration is denoted by c_1 .
2. Within each cluster — the total number of events in a cluster is divided by the total size of the cluster (using the same scale that was used for the calculation of the study area) to yield a cluster concentration denoted by c_2 .

The cluster concentration is then divided by the study area concentration to yield a value $C = c_2/c_1$. If C is greater than $1.0 + \alpha$ (where $\alpha > 0$ may be defined, as needed, to highlight those hot spots that are distinctly different from the average density in the study area), the cluster has a higher relative concentration and is considered to be a hot spot. Actionable hot spots with higher relative concentration values are therefore considered to be “hotter” than hot spots with lower relative concentration values. Within the community of analysts studying hot spots, there is an ongoing debate (see Levine, 2008) over the appropriate way to express the size of the study area. The analyst may use one of various catchment approaches — convex hull, an enclosing ellipse, a rectangle bounding all observations, or the actual jurisdiction in which the resources are to be deployed — that he/she determines to be most appropriate, but the size of the cluster in the actionable hot spot method is not subject to debate. Whether disorder events occur in all parts of an actionable cluster is irrelevant — the effective size of the cluster must reflect the resource to be deployed, not only the area in which events have historically occurred. We make the following recommendation: **When computing the concentration of events in**

a spatially constrained cluster to determine whether it meets the criteria for establishment as a valid actionable hot spot, the size of the area must reflect the entire coverage area of the resource to be deployed.

Prioritization

Once actionable clusters have been determined to be actionable hot spots, their total number may exceed the resource quantity constraints of the decisionmaker. Therefore, the actionable hot spots that are candidates for deploying resources must be prioritized in some fashion. Any reasonable prioritization must consider how effective the resource would be if deployed. Since the purpose of prioritization is to match the spatially constrained resources available in limited quantities with the problem, the prioritization should reflect the objective of the resource deployment and — if relevant to achieving that objective — the temporal constraints. For example, if the objective is to reduce burglary in a small area and the deployable resource is a police patrol car available during the midnight – 8am shift, it would make little sense to put emphasis on historical events that occur during times when the patrol car is not active. A prioritization approach should put more emphasis on disorder events that occur at roughly the same time as the resource is deployed.

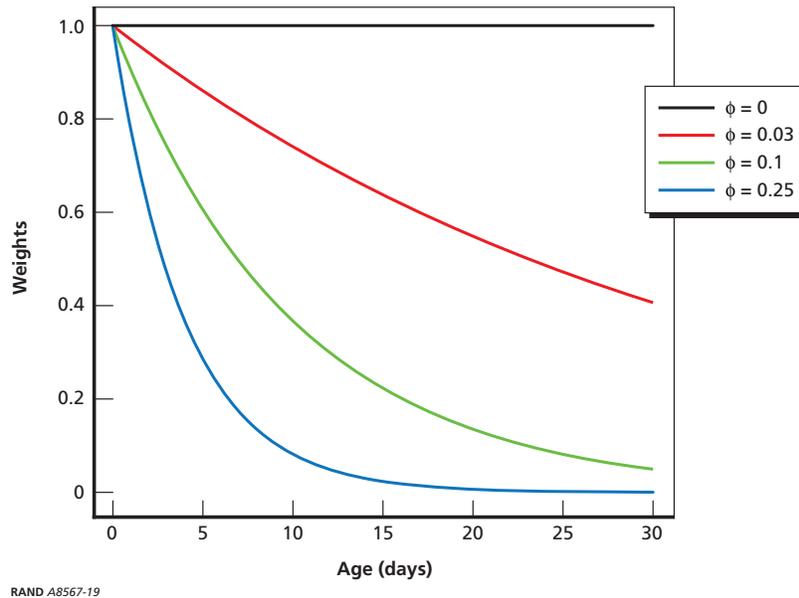
For a given objective function and known constraints, this report proposes that each candidate actionable hot spot be weighted according to how well it is synchronized with the anticipated deployment of resources meant to combat future disorder events. Additionally, the weighting process may be modified to put more emphasis on recent activity and less on historical events that have occurred in the more distant past. In the previous chapter, we gave the example of how the spatial attack patterns of piracy in the Gulf of Aden and coastal Somalia shifted over time. To account for that outcome, we recommended that older events be down-weighted or removed from the analysis. One possible family of weighting schemes that can apply differential weights to events based on their age is the simple exponential function. If w_i is the measure of importance of the i^{th} observation and A_i is the age of that observation (the time difference between the expected initiation of resource deployment and the occurrence of the historical event), the temporally weighted measure of importance is

$$w_{it} = e^{-\phi A_i}, \quad \text{where } \phi, A_i \geq 0.$$

Some discount functions that result from various values of ϕ are shown in Figure 4.1; note that for $\phi = 0$, every observation receives a weight of 1.0 and is equivalent

to not weighting at all.

Figure 4.1
Temporal Discount Functions

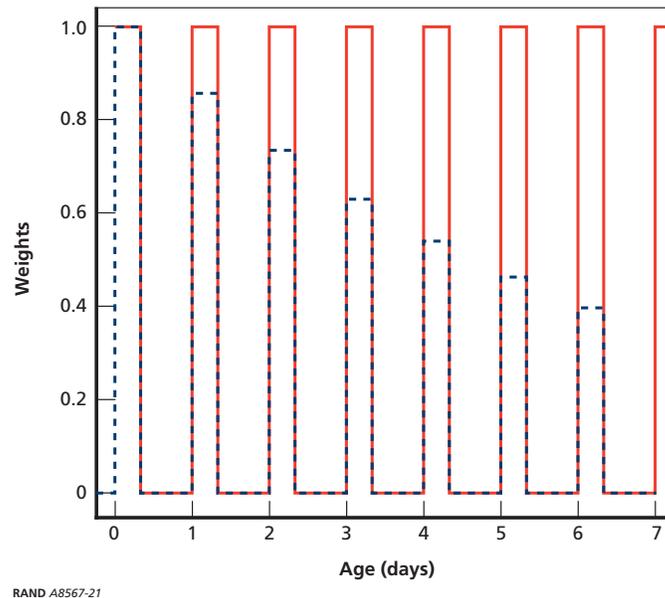


The actual value of ϕ reflecting the degree to which the cluster of the disorder event moves over time does not need to be established by the analyst. Rather, it can be found through an experimentation process that will be discussed later in this chapter.

The synchronization with the expected time of resource deployment can also be found through experimentation, but the basic shape of the weighting function should reflect knowledge of the deployment patterns. Two examples of how a weighting function may be constructed are shown in Figures 4.2 and 4.3. In Figure 4.2, the deployment of resources is expected to be in eight-hour shifts. The red weighting function yields weights of 1.0 for all historical events that have occurred during the same time window as the expected resource deployment over a one-week period (the weighting may exceed seven days in this notional example). The blue dotted line reflects a similar deployment schedule but one in which the observation weights are

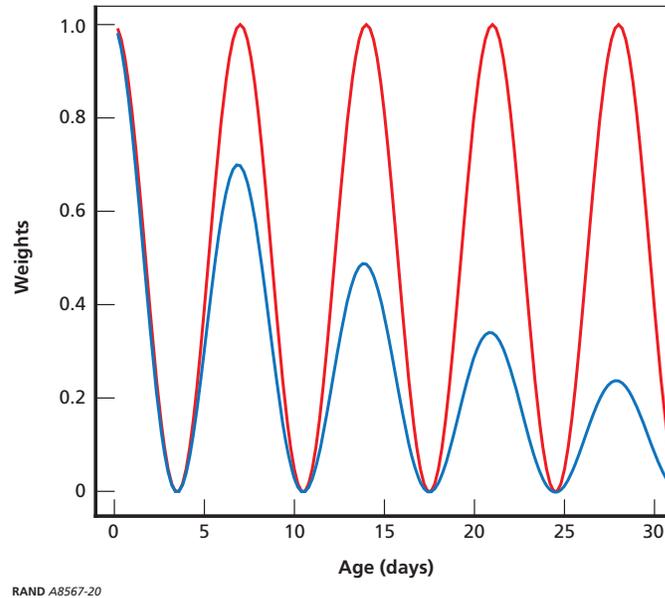
discounted over time to reflect the age of the event. Figure 4.3 shows a smoother type of weighting function (this one assumes a weekly deployment of resources, so the period of the wave is seven days). This may be more appropriate when events that fall outside of the resource deployment window are considered to have some importance but are given lesser weight than those that do fall inside the window.

Figure 4.2
Step-Function Temporal Prioritization



Weighting of observations should reflect temporal constraints to ensure synchronization with the resource under consideration for deployment. There are cases where no temporal constraints exist, but the objective function (defined in Chapter One beginning on page 4) suggests that differential weights be given to observations to ensure that the deployment of resources is correctly matched with the problem. For example, if the objective of a health care provider is to equalize compliance rates for colon cancer screening across racial/ethnic groups, the analyst could give weights to groups proportional to their historical rates of noncompliance. This would give more weight to observations in clusters populated by individuals with the lowest compliance rates and less to those areas populated by groups who have been historically

Figure 4.3
Smooth Temporal Prioritization



more compliant.

Once each observation has been appropriately weighted, a cluster score may be computed which is simply the sum of the weights in the hot spot. Prioritization then becomes simple: The actionable hot spots are ordered based on their marginal contribution to a cumulative total score (the total cumulative score will be equal to the sum of the weights for distinct events that fall within identified hot spots). Resources should then be deployed first against the actionable hot spot with the highest marginal contribution to the cumulative score, followed by the one with the second highest score, etc., until the deployable resources are depleted. Since it is possible that hot spots may overlap and so events may be counted multiple times, only distinct events (those not already included in hot spots with higher marginal contributions) are counted toward marginal cluster scores. Of course, all of the events in the highest-ranking hot spot will be used in marginal score — the process of omitting nondistinct observations need only be applied to subsequent hot spots in order to accurately

measure the marginal values. We make the following recommendation: **Among all hot spots deemed to be actionable given constraints, the ones that are best able to be addressed (because of synchronization of historical activity with expected deployment of resources and/or because they are more consistent with the objective of the introduction of resources) should be selected for resource deployment until resources are depleted.**

Measuring Expected Performance

Although it is not possible to know how effective the resource will be once deployed, it is possible to use historical data to see if — at the very least — the deployment of resources based on the application of the AHS method would have correctly selected actionable areas in which future events occur during the time window when the resource would have been deployed. The performance metric is then the total number of events that occur within the recommended actionable hot spot during the deployment period, adjusted to remove any observations that occur in the interior of more than one overlapping hot spot.

For example, if the objective is to prevent burglary by sending out patrol cars to hot spots during the midnight – 8am shift (temporal constraint), and if the cars have a patrol area of ten square city blocks (spatial constraint), and there are two patrol cars available for deployment for a period of seven days (quantity constraint), the computation of the metric would be done according to the following steps:

1. If the time when the resource deployment will begin is represented by t , weighted actionable hot spots (given the constraints) would be computed using all data available prior to t (perhaps a little earlier to account for the time to analyze historical data and prepare the resources for deployment).
2. The actionable hot spot with the highest weighted marginal scores would be selected for action, followed by the hot spot with the second-highest weighted marginal score (after removing events that are also counted in the highest-weighted hot spot).
3. For the next seven days beginning at time t , the number of burglary events that occur within each hot spot during the midnight – 8am period is counted — this is the expected performance metric. Should an event occur within more than one nominated hot spot (this occurs when hot spots overlap), each distinct event is only counted once.

With that metric, it is now possible to see if the selection of actionable hot spots was successful and by how much. This approach will allow decisionmakers comparing alternative resources for deployment to test the number of future events they would have been in the position to deter, disrupt, or prevent for each deployable asset type. Therefore, the policymaker can test **the AHS performance metric on historical data to yield an expected level of effectiveness and choose the deployable resources that would have been most effective if deployed.**

The major drawback of this approach is that, since the intervention is not actually observed, the metric does not reflect actual effectiveness but rather *expected effectiveness*. However, since the future cannot be observed until it happens, this is a reasonable metric for estimating effectiveness and comparing the potential impact that alternative resources might return. One other drawback that is not unique to the AHS method is that, in an adversarial environment, the source of the disorder may adapt to actual resource deployment to prevent detection. For example, the burglars who operate in hot spots against which resources are deployed during the deployment hours may shift their temporal pattern or displace to another area not subject to predictable police patrols. Research into this action-reaction cycle is ongoing at RAND and we expect that later reports on the AHS approach will reflect this phenomenon.

Calibration

One of the advantages of the AHS methodology that we developed is the extent to which the general approach can be tailored to a specific area and set of circumstances. The spatial, temporal, and quantity constraints on the deployable resources are known, but two other input parameters that yield larger expected effectiveness in hot spots may be found through experimentation:

- The minimum number of observations required to be considered a hot spot. Although the default value in existing geospatial software is 5, this value may not be appropriate for the phenomena being studied or the local area in which it is used to identify actionable hot spots. For example, once some criminals decide to create disorder in a selected neighborhood, they may hit ten times before moving on to another area. This does not imply that the appropriate number of observations required be ten, rather it suggests that smaller numbers be used; if the criminal strikes occur in groups of ten, there is no reason to wait very long before deploying resources to intercept him. After all, he has

historically shown that he is likely to commit eight more crimes in the same neighborhood.

- The temporal discount function, ϕ . Activity that is persistent in certain areas would be given higher weights if the value of ϕ were small or zero (note that for $\phi = 0$, all observation weights are 1.0 and so the “score” is simply the number of events in that hot spot). Higher values of ϕ are more appropriate when the spatial clustering of events tends to move over time. Note also that the number of days of historical behavior that should be included in the analysis can be reflected in the discount function by simply assigning zero weights for observations whose age exceeds a certain threshold. However, it would be computationally more efficient to simply remove those observations from the data set.

With the historical data as a training set, various combinations of these two parameters can be used to determine which combination returns the highest value of the performance metric. This evaluation with the training set of historical data serves three purposes: (1) It tunes the parameters to be as relevant as possible in the particular geographic region where this methodology is being deployed, (2) it provides a deeper understanding of the spatial and temporal patterns, (3) through this tuning, the algorithm can detect changes in environmental conditions. Once the combination of input parameters that returns the highest performance metric has been determined, it should be used to identify actionable hot spots. For a given set of constrained resources, not only does the AHS approach identify and prioritize hot spots for action, but the hot spots nominated for resource deployment take into account both the temporal and spatial dynamics of the historical disorder events.

Cross-Validation

As in most data analysis exercises, it is important that the resulting parameter estimates and hot spot nominations are not the result of *overfitting* the data to the historical data set. Careful cross-validation withholds part of the historical data during the parameter calibration phase of the analysis and then evaluates the performance of the suggested resource allocation against those events that are not part of the data that were used in the calibration. Since the spatial relationships in the data are fundamental to the generation of constrained hot spots, our proposed calibration used a temporal cross-validation. The cross-validation approach used in our

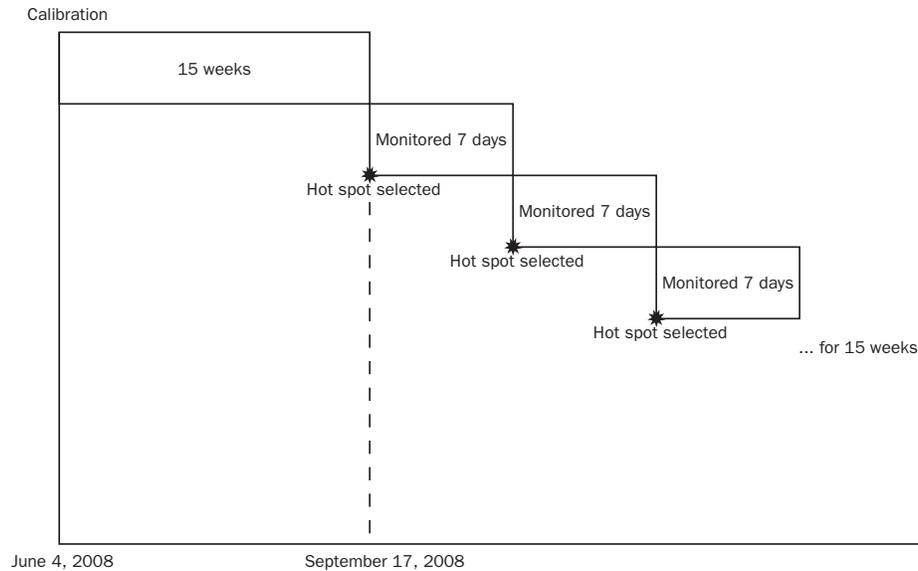
calibration process (discussed earlier) can be more formally expressed in the following steps:

1. For data indexed sequentially by a time index ($t = 1, 2, \dots, T$), we first partition the data into approximately equal halves; assuming T is an even number of time steps, the partitions are $t_1 = 1, 2, \dots, T/2$ and $t_2 = T/2 + 1, T/2 + 2, \dots, T$. With this partitioning, t_1 represents the time window containing observations used in the *training set* and t_2 represents the time window containing observations used in the *sample set* against which performance was measured.
2. The sample set, t_2 , is then divided into equal sub-partitions representing the time length in which a set of nominated hot spots is assumed to be deployed in order to measure how many disorder events occurred within those nominated hot spots. If d is the number of days in which the resource is expected to be deployed against a nominated hot spot, then monitoring occurs for a total of $p = T/2d$ deployment periods in the training set, where each of the p deployment periods contains d time steps. Using the calibration parameters that were generated in the training phase with the data in time window t_1 , a new set of hot spots is nominated for every d time step within t_2 .
3. For the hot spots nominated at the end of the $k \cdot d^{th}$ ($k = 0, \dots, p - 1$) time step following $T/2$, the number of events that fall interior to the nominated hot spot(s) in the deployment period — that is, within the next d time steps — is counted and used to compute the performance metric.

As an example of how this cross-validation is performed in the maritime piracy case study (to be explained in greater detail in Chapter Five), refer to Figure 4.4.

The historical data contained 30 weeks = 210 days of historical piracy events in the Gulf of Aden. The training set was based on the first 15 weeks ($t_1 =$ June 4, 2008 – September 16, 2008) and the sample set contained the second 15 weeks of data ($t_2 =$ September 17, 2008 – December 31, 2008). A set of input parameters (temporal discount function and minimum set of observations required to establish a hot spot) was set for the training set, t_1 , and used to nominate an actionable hot spot (the quantity constraint is assumed to be equal to one resource) on September 17, 2008. During the next $d = 7$ days, the nominated hot spot was monitored to see how many piracy events occurred within that nominated hot spot. Seven days later (on September 24, 2008), another hot spot was nominated using the data from June 4, 2008 through September 23, 2008 and that hot spot was monitored for the next $d = 7$

Figure 4.4
Cross-Validation and Effectiveness Measurement Time Line



days. Overall, this process was repeated $p = 15$ times — resulting in a total of 105 days ($7 \text{ days} \times 15 \text{ periods}$) of monitoring. The performance metric is then the number of piracy events that occurred within the distinct hot spots *only* during the seven-day period in which the individual hot spots were assumed to have resources available to intercept piracy events. This process is then repeated using several different sets of input parameters to determine which set yields the best performance. The best set of parameters from this analysis would then be used to nominate hot spots during an actual resource deployment decision process.

Summary

This chapter has demonstrated that our methodology can be used to identify hot spots against which resources can be deployed. It also provided a method for determining which hot spots are “hotter” and a method for prioritizing hot spots for resource deployment based on their expected synchronized asset deployment and the objective of the intervention. Finally, it provided a performance metric that can be used to

compare the expected effectiveness of different types of resources; the performance metric, itself, can be calibrated to reflect the temporal and spatial dynamics of the historical disorder events.

Case Studies

The actionable hot spot methodology described in earlier chapters was originally developed to help fight the IED problem in Iraq. Existing spatial analysis tools were modified to allow decisionmakers to limit the number of candidate IED hot spots to areas that conformed to the physical limits of the resources tactical commanders intended to deploy against IED emplacements. Through examples across different research areas, this chapter serves as a response to the third research question: *Can the actionable hot spots methodology be applied to guide resource allocation in research areas beyond the IED application for which it was originally developed?*

It is the authors' hope that this generalized version of the AHS methodology will find usefulness beyond the counter-IED application for which it was developed. Any decisionmaker who is faced with deploying scarce resources to geographic areas where certain types of undesirable activity or phenomena occur may find this approach useful. This approach is not intended to replace any existing spatial analysis tools but rather to augment them with the ability to conduct analysis where known constraints exist. To demonstrate the diversity of public policy areas under which this approach may be used, this report also provides three example applications; one in the maritime domain with national security implications (piracy in the Gulf of Aden), one in domestic health care delivery (colon cancer screening in a western U.S. state), and one in criminal justice (crime in a major metropolitan area). In future research, we plan to further explore the AHS approach using simulated data to determine its appropriateness given a wider variety of temporal-spatial patterns, resource constraints, and adaptation that may result as actors attempt to move their disorder events to areas that are underresourced.

We recognize that numerous models addressing resource allocation have been specified for problems related to police, fire, emergency medical services, health care, etc., in addition to the IED emplacement problem. Our case studies explore research

topics in which RAND is currently involved and where both the problem objectives and constraints have been clearly established by subject matter experts. Although solutions to domestic health care delivery can be easily handled by well-known approaches, such as the Maximal Covering Location Problem (MCLP), we believe that the AHS approach provides an alternative solution that leverages commonly used hot spot identification tools and may appeal to geospatial analysts and policymakers unfamiliar with integer-programming approaches. Our approach may also add value to those types of resource allocation problems discussed in the other case studies — and perhaps to additional topic areas as well. The prioritization phase of the AHS methodology captures shifts in spatial patterns that may occur as new target opportunities arise and/or the deployment of resources intended to interrupt future disorder events causes the actors to avoid detection. In that sense, we see the AHS approach as one possible way to address resource allocation problems when there is a repeating action-reaction exchange between those actors who deploy resources against disorder activities and those who are responsible for them.

Maritime Piracy

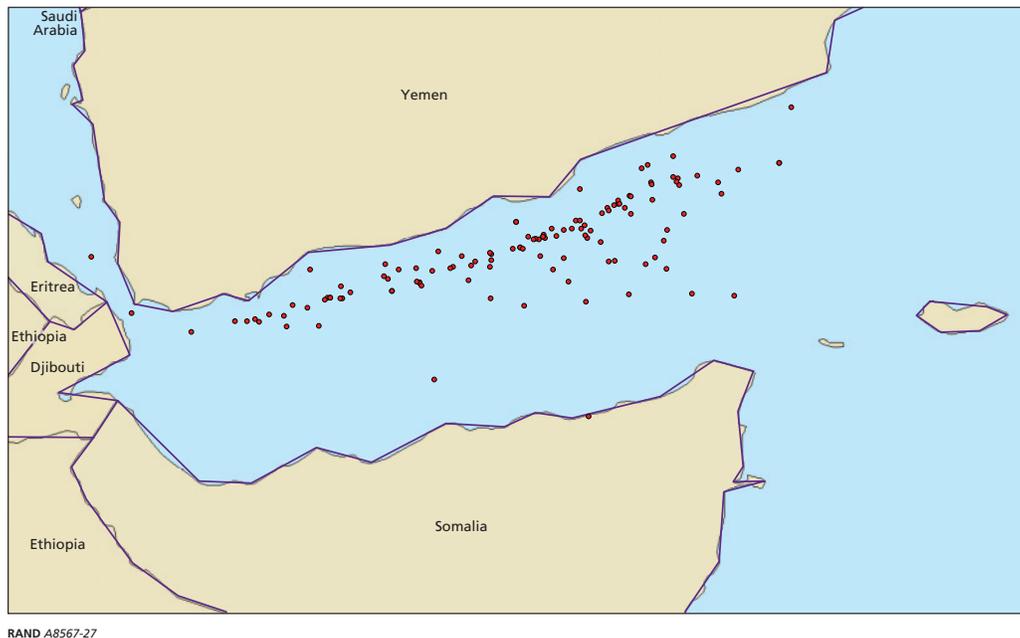
Maritime piracy is a centuries-old problem that threatens international shipping. In recent years, there has been an increasing number of pirate attacks off the coast of Somalia and in the Gulf of Aden (see Figure 5.1). In August 2008, the multinational Combined Task Force 150 established a Maritime Security Patrol Area (MSPA) in the Gulf of Aden to combat piracy in that region. In November 2008, United Nations Security Council Resolution 1838 granted nations with armed vessels the authority to exercise force to repress pirate acts. This issue moved to the forefront of U.S. security concerns when a vessel with a U.S. flag, the *Maersk Alabama*, was seized by four Somali pirates about 280 miles southeast of the port city, Eyl. Ultimately, U.S. Navy SEAL snipers killed three of the pirates and took a fourth into custody. With U.S. Naval Forces now patrolling the Gulf of Aden, a tool such as the actionable hot spots methodology might be used to find small areas preferred by pirates to conduct their attacks and launch direct counter-piracy actions given their available resources. Since the MSPA encompasses a huge geographic area, identification of actionable hot spots would enable naval commanders to focus their resources on areas that have historically demonstrated clustering. For this example, a patrolling U.S. destroyer is the deployable resource. A March 25, 2009, interview with a former Navy SEAL, Richard J. Hoffmann, yielded information about the range of the destroyer and its

operational planning cycle supporting counter-piracy efforts. With a suggested time of 40 minutes between realization that a pirate attack was imminent and the conclusion of the attack, we estimated that the single destroyer must be within 20 nautical miles in order to respond before the attack is over and the perpetrators have escaped.

Objective:	To deter, disrupt, or prevent pirate activity in the MSPA
Deployable Resource(s):	U.S. Navy destroyer
Constraints:	
<i>Spatial</i> —	The actionable response radius is 20 nautical miles (nm)
<i>Temporal</i> —	The destroyer remains afloat 24 hours per day, during which time it is available to respond to distress calls
<i>Quantity</i> —	There is only one destroyer

In this case study, we assumed that the naval task group selects a position to locate its destroyer at the beginning of each week based on recent piracy attack patterns. Every seven days, the patterns are reanalyzed and the destroyer moves to the position where it is expected to be within range of the most future piracy attacks that are expected to occur. We calibrated input parameters experimentally to determine (1) the number of days of historical events that should be used to determine an actionable hot spot, (2) whether the events should be weighted over time to put less emphasis on observations occurring in the window, and (3) the minimum number of observations required before a cluster is eligible to become a hot spot. The data were calibrated on a weekly basis for the 15 weeks prior to September 17, 2008, and a single actionable hot spot was selected weekly based on a prioritization of all candidate hot spots that were within the resource constraints. The highest performance level was returned when the 45 days of historical data were used (selected from 15, 45, or 90 days), when equal weight was given to each observation (compared with a lesser weight for older observations), indicating persistence in the location of hot spots, and when the minimum cluster size was 3 (selected from 2, 3, 4 or 5). We tested all possible combinations, and the combination of these three input parameters yielded an expected effectiveness of 40 events in 105 days. This means that 40 piracy events occurred within the 20 nautical miles (nm) of the destroyer's position within a week after the position was selected. Over the 15-week period, a total of 15 hot spots were identified — one per week — and a new position was located each week based on

Figure 5.1
Piracy Incidents in the Gulf of Aden, July – December 2008



Source: National Geospatial Intelligence Agency.

historical data available prior to the positioning recommendation.

Given the parameters that yielded the maximum effectiveness during the calibration phase, we used the data from September 17, 2009, to December 31, 2009, to measure performance. Overall, 15 actionable hot spots were identified (one per week using only the last 45 days of data to select and prioritize actionable hot spots). One or more piracy events occurred within the identified hot spots during eight of the 15 different weeks in the evaluation period. In total, there were ten piracy events within the actionable hot spots during the 15 weeks. There was less piracy activity in the study area than during the calibration, but, on one day, there were five distinct clusters of observations that met the criteria for resource deployment (having three or more events within a radius of 20 nautical miles within the last 45 days). This indicates that the AHS approach can be useful in selecting and prioritizing actionable clusters when the data indicate several areas that should be considered for resource deployment. The success rate is much lower than the other case studies, but the

size of the area is an important consideration: During the 15-week evaluation period, there were only 131 reported piracy incidents over the entire Gulf of Aden/Somali coast. Given that the total area is thousands of square kilometers and there was a small number of events in this time period, we found the ability to allocate resources against ten of those events in a circle with a radius of 20 nm to be fairly encouraging since only one destroyer was assumed to be capable of deploying against piracy over a huge area.

Colon Cancer

RAND Health's research portfolio includes various projects supporting efforts of health care providers in different regions to improve the quality and outcomes of care among the diverse populations they serve. For example, a major concern of health care plans and providers located in the western United States is the relatively high rate and poor outcomes of colon cancer among large minority groups, such as Hispanics. Such gaps are thought to be due, in part, to very low rates of preventive screening (e.g., colonoscopy) among certain minority groups. That is, even when insured, certain minority groups tend to have substantially lower rates of screening than comparable white patients. Thus, many health plans and providers are interested in finding more efficient and effective ways to identify where groups of high risk members live and target interventions to increase colon cancer screening.

In the case example that follows, a group of health providers based in the western United States were interested in (a) establishing screening clinics and public outreach information campaigns in local neighborhoods of the county in the study area with high levels of noncompliance among minorities and (b) staffing a clinic with personnel who understand the languages and cultural sensitivities in the selected neighborhoods. For confidentiality purposes, the identity of the health care providers and the participants in the particular county considered will remain anonymous. In addition, the specific geographic location has been masked by shifting the longitude by a fixed amount, selecting a rectangular region that encloses the data, and removing the underlying data layers depicting roads, streams, and place names that would allow the area to be identified. Nevertheless, the case represents a subset of actual subscriber data that exhibits real instances of noncompliance with recommendations for colon cancer screening. It should be noted here that the MCLP would be an entirely appropriate — and perhaps more efficient — way of addressing this problem, since the disorder events are static and there is unlikely to be any attempts by the unscreened

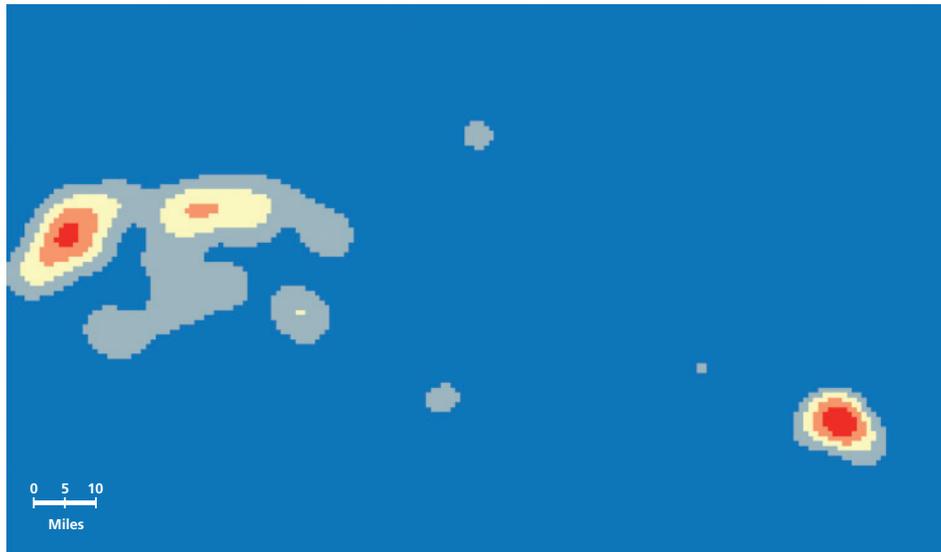
subscribers to change address so they fall outside of the reach of established clinics.

Objective:	To increase colon cancer screening compliance among minorities in the subscriber database
Deployable Resource(s):	Screening clinics and public outreach information campaigns staffed with culturally sensitive personnel
Constraints:	
<i>Spatial</i> —	The actionable response radius is either 1 or 2 km — within walking distance of the clinic and within the capacity limits of a clinic capable of handling colon cancer screening in a dense, urban area
<i>Temporal</i> —	None
<i>Quantity</i> —	There is funding for two small clinics or one large one

The data set of subscribers contains 1,753 observations, of which a minority has not complied with health care providers' recommendations for colon cancer screening as of December 2006. For the analysis of alternatives, the provider is comparing the relative potential effectiveness that would result from setting up two small clinics augmented with outreach information campaigns that have a radius of 1 km or one large clinic/outreach effort with a radius of 2 km. An unconstrained hot spot analysis using the kernel density approach indicates two potential areas for locating the screening clinics (see Figure 5.2). This example provides a good example of why constraints need to be considered during hot spot analysis:

- The size of the two hot spots generated by the KDE approach exceeds the spatial constraint (the red area in the lower right is approximately 2–4 times larger than the actionable response radii) so it is unclear which sub-region within the larger hot spots should be targeted for intervention.
- It is unclear which of the nonactionable hot spots is “hotter” and should be selected for resource deployment in order to yield the largest effect on the minority colon cancer rate.
- The KDE approach does not allow a common baseline against which the effectiveness of alternative courses of action can be measured.

Figure 5.2
KDE of Minority Colon Cancer Screening Noncompliance



RAND A8567-24

Source: Anonymous Health Care Provider/RAND Corporation.

The highest-prioritized hot spots were the ones with the greatest amount of minorities noncompliant with color cancer screening within the radius of the clinic. For the alternative with two small clinics, the largest numbers of noncompliant individuals were 80 and 64, respectively, yielding a total impact of 144 individuals (8.2 percent of the entire target population) that may be targeted for outreach. With the single, large clinic alternative — the largest number of individuals within a circle of radius 2 km is only 103 individuals (5.9 percent of the entire target population). This case study demonstrates the benefits of using the AHS approach and indicates its potential usefulness in the health care delivery field. The AHS approach allows a comparison of the effectiveness of alternatives based on an assumed deployment of constrained resources.

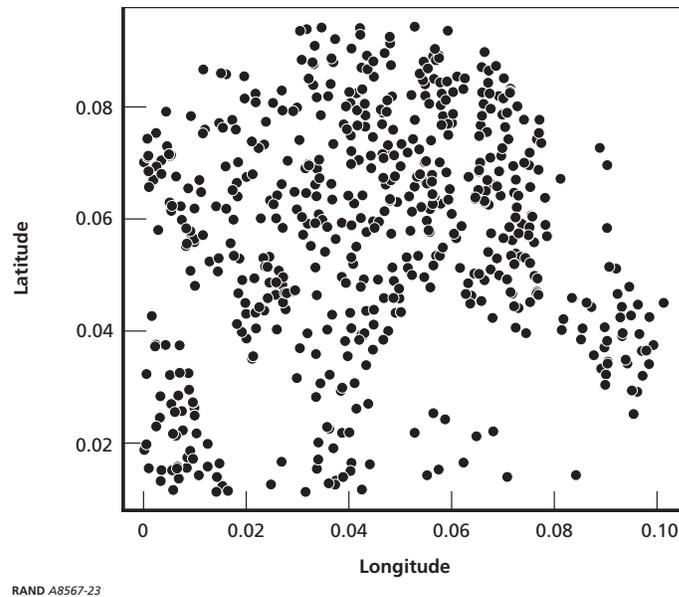
Metropolitan Crime

A major issue facing a metropolitan area in the southern United States is the high burglary rate that has plagued the city. The city's police department is divided into several distinct precincts, each one allocated funding to address the most pressing crime problems in that area. Operations commanders in each precinct must allocate scarce resources with the knowledge that an increase in policing resources in one area will result in less attention to crime problems elsewhere in their jurisdiction. In this case study, the commander of Precinct 1 must decide where to position patrol cars to create the largest possible impact in deterring, disrupting, or preventing burglary in the area. For confidentiality purposes, the identity of the metropolitan area and its criminal activity will remain anonymous. The data presented in this case study have been masked by shifting the longitude of historical events by a fixed amount, selecting a rectangular region of the city to obscure its shape, and removing the underlying data layers that would allow the metropolitan area to be identified. However, the study area represents a subset of actual crime data that exhibits real instances of burglary from December 16, 2008, to July 15, 2009. The available police patrol times are midnight – 8am, 8am – 4pm, and 4pm – midnight. The historical data containing 1,261 historical burglary incidents indicate that 52 percent of burglaries in the precinct during the analysis period occurred during the 8am – 4pm shift, with 41 percent occurring in the 4pm – midnight shift and only 7 percent occurring between midnight and 8am. For that reason, the additional patrol car will be deployed to a hot spot during the 8am – 4pm shift. The 658 observations occurring during the 8am – 4pm shift are shown in Figure 5.3. The density of observations indicates how difficult it may be to choose an actionable hot spot from this 19-square-mile precinct.

Objective:	To deter, disrupt, or prevent burglary in Precinct 1 using patrol cars
Deployable Resource(s):	Patrol cars
Constraints:	
<i>Spatial</i> —	A patrol car may serve a 10-square-block area (1 mi x 1 mi)
<i>Temporal</i> —	The patrol car is available from 8am to 4pm daily
<i>Quantity</i> —	Funding allows for 1 patrol car

Only historical burglary events that occurred within the 8am – 4pm shift were included in the analysis. Calibration of input parameters was conducted to experi-

Figure 5.3
Daytime Burglary Events, December 16, 2008, to July 15, 2009



Source: Anonymous Metropolitan Police Department.

mentally determine (1) the number of days of historical events that should be used to determine an actionable hot spot, (2) whether the events should be weighted over time to put less emphasis on observations that occurred earlier in the window, and (3) the minimum number of observations required before a cluster was eligible to become a hot spot. The data were calibrated on a daily basis for the 15 weeks prior to April 1, 2009, and a single actionable hot spot was selected daily based on prioritization of all candidate hot spots that were within the resource constraints. The greatest number of events occurred when 28 days of historical data were used (selected from 7, 14, 28, or 35 days); equal weight was given to each observation (compared with a lesser weight for older observations), indicating persistence in the location of hot spots; and the minimum cluster size was 5 (selected from 2, 3, 4, or 5). All possible combinations were tested, and the combination of these three input parameters yielded an expected effectiveness of 32 percent. This means that $34 \text{ } 34 / 105 = 32$ percent) burglaries occurred within the ten-square-block patrol area selected as the most actionable hot

spot during the 8am – 4pm window over the 15-week period (a total of 105 hot spots were identified — one per day).

With the parameters that yielded the maximum effectiveness during the calibration phase, we used the data from April 1, 2009, to July 15, 2009, to measure expected performance. Overall, 105 actionable hot spots were identified (one per day using only the last 28 days of data to select and prioritize actionable hot spots). One or more burglary events occurred within the identified hot spots during the 8am – 4pm shift on 35 different days. In total, 44 burglaries occurred within the actionable hot spots during the 105 days, which gives an expected effectiveness of $(44/105 = 42 \text{ percent})$. There was significant burglary activity in the study area, and on one day there were nine distinct clusters of observations that met the criteria for resource deployment. This indicates that the AHS approach can be useful in selecting and prioritizing actionable clusters when the data indicate several areas that should be considered for resource deployment. In addition to helping choose the “hottest” hot spots that are within resource constraints (and very small in this example), this case study also indicates that the AHS approach has potential beyond the counter-IED application for which it was originally developed. Selection of a new one-square-mile area to patrol daily in a 19-square-mile area based on historical data yielded a relatively high success rate of 42 percent, indicating that the AHS approach may be useful to law enforcement personnel.

Summary

This chapter has demonstrated that the AHS methodology created for a specific application (counter-IED operations) can be applied in other research areas. Augmentation of existing analyses with the AHS methodology will allow geospatial analysts not only to conduct hot spot analysis using their standard toolkit but to be able to do so while considering resource constraints. The approach also allows policymakers to compare alternative suites of resources to determine which is expected to generate a larger impact on reducing the disorder events that they are attempting to deter, disrupt, or prevent. The success of the approach is based on the degree to which clustering is present in the data and the ability to deploy available resources that can be spatial and temporally matched against the disorder activity.

Implications

Decisionmakers tasked with deterring, interrupting, or preventing undesired activities are limited by constraints caused by available, scarce resources; often these resources cannot cover the vast geographic areas in which the problems occur. In the extensive body of research addressing the use of spatial analysis in criminal analysis, pattern recognition of insurgent and terrorist activity, and public health, the term *hot spot* has been adopted to indicate areas where there exists a greater than average number of problem events. This technical report provides a methodology that can be used to select and prioritize hot spots against which constrained resources can be deployed. The methodology provides a means of measuring the expected effectiveness that would result by deploying resources against a problem using scarce resources. Not only does this approach provide a tool for aiding the decisionmaker as he/she chooses how to allocate existing resources, it also provides a mechanism for comparing the potential effectiveness of alternative resources.

The “actionable hot spot” methodology is not intended to replace any of the existing tools widely used by spatial analysts. Rather, it provides an enhancement to hot spot detection algorithms by enabling geospatial analysts to match problem areas with the resources that they plan to deploy to combat the underlying problem. Users of *CrimeStat*[®], *GeoData*TM, and *ArcGIS*[®] across many fields may find utility in this approach when they are faced with constrained resources. Originally developed for a particular application, combating IED emplacement in Iraq, the approach had obvious applications in other fields. By modifying the original application to make it generalizable across a broad array of research topics, we have created a policy decision tool that may find utility across many areas (see Table 6.1 for a nonexhaustive list of potential applications).

Table 6.1
Potential Applications of Actionable Hot Spot Methodology

Topic	Application	Deployable Resource
National security	Maritime piracy	Visual surveillance assets Armed surface ships
	Counter-IED/indirect fire	Snipers Visual surveillance assets Infrared detectors Quick reaction forces
	Insurgent network detection	Visual surveillance assets Signal direction-finding assets
Homeland security	Border integrity	Visual surveillance assets Acoustic surveillance assets Border patrol agents
Criminal justice	Law enforcement	Police patrols Visual surveillance assets Task forces
Health	Disease prevention	Screening clinics Targeted public service campaigns
	Pandemic crises	Immunization clinics Targeted public service campaigns
Labor and population	Economic disparity	Employment programs Poverty assistance

Bibliography

Benigni, Matthew, and Reinhard Furrer, "Periodic Spatio-Temporal Improvised Explosive Device Attack Pattern Analysis," Technical Report MCS-04-08, Golden, CO: Colorado School of Mines, Department of Mathematical and Computer Sciences, 2008.

Bowers, Kate J., Shane D. Johnson, and Ken Pease, "Prospective Hot-Spotting: The Future of Crime Mapping?" *British Journal of Criminology*, Vol. 44, 2005, pp. 641 - 658.

Braga, Anthony A., "Hot Spots Policing and Crime Prevention: A Systematic Review of Randomized Controlled Trials," *Journal of Experimental Criminology*, Vol. 1, 2005, pp. 317 - 342.

Burch, James W., and Michael N. Geraci, "Data-Driven Approaches to Crime and Traffic Safety," *Police Chief*, July 2009, pp. 18 - 23.

Cameron, James G., and Michael Leitner, "Spatial Analysis Tools for Identifying Hot Spots," in *Mapping Crime: Understanding Hot Spots*, National Institute of Justice, Technical Report NCJ 20939, 2005.

Chainey, Spencer P., "Methods and Techniques for Understanding Crime Hot Spots," in *Mapping Crime: Understanding Hot Spots*, National Institute of Justice, Technical Report NCJ 20939, 2005.

Chainey, Spencer P., S. Reid, and N. Stuart, "When Is a Hotspot a HotSpot? A

Procedure for Creating Statistically Robust Hotspots of Crime,” in *Innovations in GIS 9*, London: Francis and Taylor, 2002.

Chetwynd, A. G., and P. J. Diggle, “Second Order Analysis of Space-Time Clustering,” *Statistical Methods in Medical Research*, Vol. 4, 1995, pp. 124 - 136.

Church, R. L., “The Planar Maximal Covering Location Problem,” *Journal of Regional Science*, Vol. 24, 1984, pp. 185 - 201.

Church, R. L., and C. S. ReVelle, “The Maximal Covering Location Problem,” *Papers of the Regional Science Association*, Vol. 32, 1974, pp. 101 - 118.

Duda, Richard O., and Paul E. Hart, *Pattern Classification and Scene Analysis*, New York: John Wiley and Sons, Inc., 1973.

Eck, John E., “Crime Hot Spots: What They Are, Why We Have Them, and How to Map Them,” in *Mapping Crime: Understanding Hot Spots*, National Institute of Justice, Technical Report NCJ 209393, 2005.

Eck, John E., Spencer P. Chainey, James G. Anderson, Michael Leitner, and Ronald E. Wilson, *Mapping Crime: Understanding Hot Spots*, National Institute of Justice, Technical Report NCJ 209393, 2005.

Ge, Rong, Martin Ester, Wen Jin, and Ian Davidson, “Constraint-Driven Clustering,” *The 13th International Conference on Data Mining and Knowledge Discovery*, San Jose, CA, 2007.

Gentle, James E., *Computational Statistics*, New York: Springer-Verlag, 2002.

Gesler, W. M., and D. P. Albert, “How Spatial Analysis Can Be Used in Medical Geography,” in D. P. Albert, W. M. Gesler, and B. Levergood (eds.), *Spatial Analysis, GIS, and Remote Sensing Applications in the Health Sciences*, Chelsea, MI: Ann Arbor Press, 2000.

Gordon, A. D., *Classification*, 2nd ed., Boca Raton, FL: Chapman and Hall/CRC, 1999.

Gorr, Wilpen, and Andreas Olligschlaeger, "Crime Hot-Spot Forecasting: Modeling and Comparative Evaluation," U.S. Department of Justice, Technical Report 195168 (Award 98-IJ-CX-K005), 2002.

Groff, Elizabeth R., and Nancy G. La Vigne, "Forecasting the Future of Predictive Crime Mapping," *Crime Prevention Studies*, Vol. 13, 2002, pp. 29 - 57.

Haining, R., *Spatial Data Analysis: Theory and Practice*, New York: Cambridge University Press, 2003.

Heath, Garrett D., Ryan Keefe, and Thomas Sullivan, "Identifying Insurgent Indirect Fire Points of Origin for Targeting," *The 75th Symposium of the Military Operations Research Society*, Annapolis, MD, 2007.

Hofman, Sgt. Steven, "Civilian Sector Developing Counter IED Tools," *The Maple Leaf*, Vol. 10, No. 22, p. 10, 2005.

Hollywood, John S., Walter Perry, and Thomas J. Sullivan, "Targeting IED Networks in Iraq," Santa Monica, CA: RAND, 2009. Not releasable to the general public.

Kaufman, L., and P. J. Rousseeuw, *Finding Groups in Data: An Introduction to Cluster Analysis*, Wiley Series in Probability and Mathematical Sciences, New York: John Wiley and Sons, Inc., 1990.

Knox, G., "Detection of Low Density Epidemicity," *British Journal of Preventative and Social Medicine*, Vol. 18, 1963, pp. 121 - 127.

Kopka, Helmut, and Patrick W. Daly, *A Guide to L^AT_EX*, Reading, MA: Addison-Wesley, 1999.

Levine, Ned, *Crime Stat 1.0, A Spatial Statistics Program for the Analysis of Crime Incident Locations*, Washington, DC: U.S. Department of Justice, National Institute of Justice, 1999.

Levine, Ned, "The 'Hottest' Part of a Hotspot: Comments on 'The Utility of Hotspot

Mapping for Predicting Spatial Patterns of Crime,” *Security Journal*, 2008, Vol. 21, No. 4, pp. 295 - 302.

McGuire, P. G., and D. Williamson, “Mapping Tools for Management and Accountability,” paper presented at the *Third International Crime Mapping Research Center Conference*, Orlando, FL, 1999.

Openshaw, Stan, *The Modifiable Areal Unit Problem*, Norwich, CT: Geo Books, 1984.

Page, Douglas, “Police Behaving Predictably: The Other Enemy,” *Law Enforcement Technology*, Vol. 36, No. 2, pp. 74, 76 - 78, 2009.

Sullivan, Thomas J., *Classification Methods for Augmented Arc-Weighted Graphs*, Ph.D. thesis, Fairfax, VA: George Mason University, School of Computational Sciences, 2002.

Tobler, Waldo, “A Computer Movie Simulating Urban Growth in the Detroit Region,” *Economic Geography*, Vol. 46, 1970, pp. 234 - 240.

Williamson, D., S. McLafferty, V. Goldsmith, J. Mallenkopf, and P. McGuire, “A Better Method to Smooth Crime Incident Data,” *ESRI ArcUser Magazine*, January - March 1999, pp. 1 - 5.

Wilson, Ronald E., “Conclusions,” in *Mapping Crime: Understanding Hot Spots*, National Institute of Justice, Technical Report NCJ 209393, 2005.