

The Regional Economics Applications Laboratory (REAL) is a cooperative venture between the University of Illinois and the Federal Reserve Bank of Chicago focusing on the development and use of analytical models for urban and regional economic development. The purpose of the **Discussion Papers** is to circulate intermediate and final results of this research among readers within and outside REAL. The opinions and conclusions expressed in the papers are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago, Federal Reserve Board of Governors or the University of Illinois. All requests and comments should be directed to Geoffrey J. D. Hewings, Director, Regional Economics Applications Laboratory, 607 South Matthews, Urbana, IL, 61801-3671, phone (217) 333-4740, FAX (217) 244-9339. Web page: www.uiuc.edu/unit/real

VISUALIZING SPATIAL AUTOCORRELATION WITH
DYNAMICALLY LINKED WINDOWS
by

Luc Anselin, Ibnu Syabri, Oleg Smirnov and Yanqui Ren

REAL 01-T-10 November, 2001

Visualizing Spatial Autocorrelation with Dynamically Linked Windows¹

LUC ANSELIN, IBNU SYABRI, OLEG SMIRNOV and YANQUI REN

Regional Economics Applications Laboratory (REAL)
University of Illinois, Urbana-Champaign
Urbana, IL 61801

Abstract

Several recent efforts have focused on adding exploratory data analysis functionality to geographic information systems (GIS) by integrating established statistical software with a GIS. In this paper, we outline an alternative approach, where the functionality is built from scratch, using a combination of small libraries of dedicated functions, rather than relying on the full scope of existing software suites. The suggested approach is modular and freestanding. Within an overall framework of dynamically linked windows, it combines a cartographic representation of data on a map with traditional statistical graphics, such as histograms, box plots, and scatterplots. In addition, it includes several devices to visualize spatial autocorrelation in lattice (or regional) data, such as the Moran Scatterplot and LISA maps. This new program (*DynESDA2*) also implements the capability to brush polygon coverages, simultaneous linking of multiple maps with multiple statistical graphics, and interactive LISA maps.

Key Words: spatial statistics, exploratory spatial data analysis, visualization, dynamic graphics, GIS.

1 Introduction

With the proliferation of user-friendly geographic information systems, the easy access to geocoded data and the increased interest in substantive “spatial” research questions, the demand for sophisticated spatial analytical tools has increased considerably in recent years (Goodchild *et al.* 2000). Building on an early interest in linking standard statistical software packages and GIS in order to carry out generic data analyses, several more recent efforts have focused on adding exploratory *spatial* data analysis (ESDA) functionality. The approach commonly taken is to establish a link between a statistical package and a GIS by means of remote procedure calls and a client/server architecture (often referred to as “close coupling”). There are by now quite a few implementations of this idea, for example, linking statistical software packages such as S-Plus, XGobi, or XploRe to GIS software, such as ArcView and Arc/Info (for recent reviews of the relevant literature, see Anselin 1998, 2000, Symanzik *et al.* 2000, Zhang and Griffith 2000, Wise *et al.* 2001).

In this paper, we report on an “in progress” software tools development project to facilitate exploratory spatial data analysis. We approach this problem in a different way than most recent linking efforts in that the ESDA functionality is built from scratch, rather than

¹Contact: anselin@uiuc.edu. This research was supported in part by NSF Grants SBR-9410612, and BCS-9978058, to the Center for Spatially Integrated Social Science (CSISS), and by a grant from the National Consortium on Violence Research (NCOVR). NCOVR is supported under grant SBR-9513040 from the National Science Foundation. The considerable contributions by Shuming Bao to the development of the initial versions of the SpaceStat-ArcView linkage software are gratefully acknowledged. Programming help was provided by Widodo Baroka. We also thank Suahasil Nazara for kindly providing the data and GIS from the SUSENAS Indonesian Economic Survey.

by connecting existing software suites. This is accomplished by using a combination of small libraries of dedicated functions as software “components.” The approach is modular and extensible, as well as completely freestanding. It implements distributed computing, allowing the analytical engine and the data to be stored on different machines, and facilitates the use of data formats from different vendors.

The current set of tools is built using ESRI’s MapObjects software components to implement mapping and data base access, augmented with functionality to carry out the *spatial* analysis (written in C++). This does not require the use of a particular GIS and allows for the visualization of spatial data stored in different formats. The user interface consists of fully dynamically linked windows that include multiple cartographic (thematic) representations of data on maps as well as traditional statistical graphics, such as histograms, box plots, and scatterplots. In addition, it includes several devices to visualize spatial autocorrelation in lattice (or regional) data, such as the Moran Scatterplot and LISA maps (Anselin 1995, 1996, 2000).

In the remainder of this paper, we first provide some background and situate our approach among a number of other efforts that have been reported in the recent literature. We then proceed by outlining the architecture and functionality of the software, followed by a simple empirical illustration dealing with the spatial distribution of household consumption in the greater Jakarta (Indonesia) metropolitan region, or Jabotabek. We close with some concluding comments on future directions.

2 Background

The current framework, referred to as *DynESDA2*, is our latest iteration in an ongoing effort to augment the visualization and spatial data manipulation functionality of a GIS with an analytical engine that contains spatial statistical and spatial econometric methods.² The original outline of the conceptual framework for such an integration can be found in Anselin and Getis (1992) and Goodchild *et al.* (1992). The first implementation of our integrated framework consisted of interfacing the spatial econometric and ESDA functionality of SpaceStat (Anselin 1992, 1999b) with ESRI’s Arc/Info GIS in Anselin *et al.* (1993). The interaction between the two software packages was based on so-called “loose coupling,” which consisted of moving data and location-specific results back and forth between SpaceStat as the analytical engine and Arc/Info as the visualization engine. This early effort was more a proof of concept than a practical tool, as it suffered from performance problems and limitations for dynamic interaction due to the design of Arc/Info (as well as from the use of two different operating systems, Unix for Arc/Info and DOS for SpaceStat). Arc/Info was used as the basis for an integrated or linked framework by a number of others as well (in a Unix environment), although using a different architecture. For example, dynamically linked windows from XGobi were interfaced with Arc/Info by Symanzik *et al.* (1994b), based on a client/server architecture, but only limited interaction was provided with the maps in Arc/Info. Similarly, the extension of Arc/Info with spatial statistical functionality implemented in the SAGE Project uses a client/server architecture to avoid performance problems with loose coupling (Haining *et al.* 1996, 1998, 2000, Wise *et al.* 2001).

The popularization of the ArcView desktop GIS software in the mid-1990s saw this GIS become the primary focus of extension efforts.³ Initially, spatial statistical functionality

²More extensive descriptions of the evolution of software tools can be found in Anselin *et al.* (1993), Anselin and Bao (1996, 1997), Bao and Anselin (1998), Anselin and Smirnov (1999a,b) and Anselin (2000).

³Other popular desktop GIS software, such as MapInfo, has only seen limited use as a platform to implement spatial statistical extensions. A rare example is Wall and Devine (2000).

was added by means of the built-in Avenue scripting language. For example, Zhang and Griffith (1997) provide limited spatial autocorrelation statistics through the application of Avenue scripts. Also, in Anselin and Bao (1997) and Bao and Anselin (1998), Avenue scripts are used to implement the link with SpaceStat. Performance problems (both in terms of speed as in terms of the size of problems that could be handled) quickly led to the adoption of different designs. Popular among these was the use of remote procedure calls to link ArcView with other (statistical) software. For example, this is applied in the series of integration efforts in a Unix environment between exploratory software such as XGobi and XploRe on the one hand, and the ArcView GIS on the other hand, by Cook, Symanzik and co-workers (see, e.g., Cook *et al.* 1996, 1997, Symanzik *et al.* 1994a, 1997, 1998, 2000). Similarly, the link between the S-Plus statistical software and ArcView is based on remote procedure calls (implemented in Avenue scripts), allowing S-Plus commands to be invoked from within ArcView and vice versa (Bao *et al.* 2000). While exploiting the functionality of ArcView for interactive mapping and querying, combined with the linking and brushing capabilities in the EDA software, these interfaces were still limited by the constraints on the number of links that could be open at the same time (a limitation of the remote procedure call implementations on these systems). Also, to the extent that they relied on built-in Avenue scripts for some spatial data handling functionality, they tended to be slow and limited in the number of spatial objects that could be handled.

The SpaceStat and DynESDA extensions for ArcView in a MS Windows environment (Anselin and Smirnov 1999a,b, Anselin 2000) were designed to address some of these performance issues. While they also suffer from some of the limitations imposed by the ArcView software, performance bottlenecks (particularly for intensive numerical operations) due to the use of the Avenue scripting language were avoided. Rather than using Avenue for computations, the main analytical engine for the statistical operations is contained in a number of dynamically linked libraries, written in C++. This forms the immediate precursor to the current *DynESDA2* implementation in terms of most of its statistical functionality. However, we also felt that some of the overhead associated with using a “complete” GIS could be avoided, especially for users more interested in data analysis than data manipulation. Therefore, the (rudimentary) mapping and data base access functionality in *DynESDA2* was constructed using ESRI’s MapObjects software components. These do not require ArcView or any other GIS to be open and allow the software to operate fully independently. In that respect, *DynESDA2* is similar in spirit to the various descendants of the original Spider software (Haslett *et al.* 1990, 1991), which implement dynamically linked windows in a self-contained framework (i.e., not relying on a GIS for mapping), where the “map” is but one of several linked views.⁴ In addition to being freestanding, *DynESDA2* also includes a number of other advances over its predecessors, such as the capability to brush polygon coverages, simultaneous linking of multiple maps with multiple statistical graphics, and interactive LISA maps. We consider this functionality more closely below. First, we turn to the overall design and architecture of the integrated framework.

3 Architecture and Design

The *DynESDA2* framework is conceptualized as a collection of modules that each handle a different aspect of the user interaction with the data. Four major modules can be distinguished, respectively dealing with data retrieval and data base management, thematic mapping, statistical analyses, and queries, as illustrated in Figure 1. Both data access and

⁴See Unwin (1996) and Wilhelm and Steck (1998) for recent examples. Similar ideas underlie the Tcl/Tk based *cdv* toolkit of Dykes (Dykes 1997, 1998) as well as Brundson’s exploration of local spatial association using a dynamically linked “map” constructed with tools available in *xlispstat* (Brundson 1998).

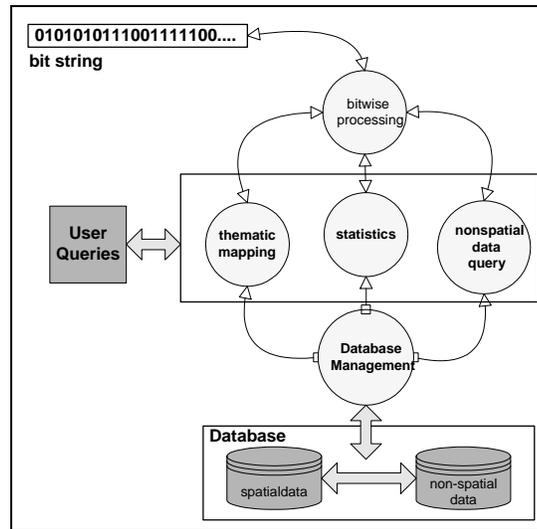


Figure 1: Basic architecture of *DynESDA2*.

thematic mapping have been implemented by customizing ESRI's MapObjects components, whereas the statistical analyses and data queries were developed in C++, using the dynamically linked libraries from the earlier version of *DynESDA* as a point of departure (see Anselin and Smirnov 1999a, Anselin 2000).

The modules are tied together in a Microsoft Windows "multiple document interface" (or MDI), with each type of analysis (mapping, descriptive statistics, spatial statistics) corresponding to a distinct type of "document." All windows are implemented using Microsoft Foundation Classes (MFC) to provide a consistent look and feel.⁵ In total, there are five generic classes of windows, each enabling a different "view" of the data: map, histogram, box plot, scatterplot and table. A specialized form of the scatterplot is used for the Moran scatterplot, and special instances of the map view yield the Moran significance map and the LISA map.

An analysis is initiated by loading a data set (such as an ESRI shapefile), which contains both the data table as well as a digital boundary file for visualization of the map. While different maps may be constructed from the same data table, the table itself is unique in each "analysis." Specific analytical functions are invoked through the Explore and Thematic Map menus, as illustrated in Figures 2 and 3. Each of the Explore menu items starts a new window, whereas the Thematic Map menu items implement a specific form of visualization for the base map (multiple such visualizations can be open at the same time).

The different views of the data are synchronized by means of a so-called bit string (or bitmap). This stores the "selected" observations in an efficient manner and is updated any time an observation or set of observations is highlighted by means of a pointer in any one of the views. Two types of selections are implemented: spatial and non-spatial. Non-spatial queries are built from SQL statements and select items from the data table that match a given set of criteria. This can also be carried out by pointing and clicking on records in the "Table" view. Spatial queries are implemented by interactively clicking or clicking and dragging on a map view, using one of the geometric shapes provided in the Select pop up

⁵In other applications that link ArcView to a statistical package, the GUI is a combination of the interfaces for each separate product, which may lead to confusion and inefficiencies. By choosing MFC as the building blocks for the interface throughout, this is avoided in the current design.

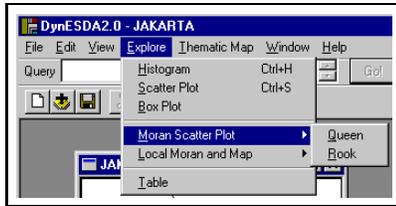


Figure 2: Explore Menu.



Figure 3: Map Menu.

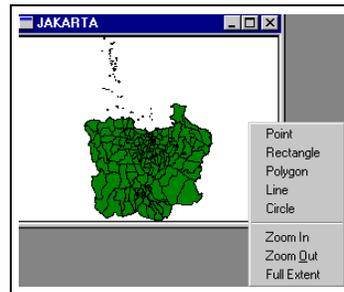


Figure 4: Select Pop Up Menu.

menu to graphically delineate the selection. Five such shapes are currently included: point, rectangle, polygon, line and circle (see Figure 4).

Dynamic linking is implemented by refreshing all windows with a new selection any time the bit string is altered. This central processing of the bit string ties all the interfaces together, as illustrated in Figure 5. The user can “enter” the tools from any number of views, such as the map, the data table or any of the statistical graphs. Only a single data table can be active at one time, since it defines the available variables and the locations of the observations, but the number of maps and graphs linked together is unrestricted. This removes a limitation that is present in many other current implementations, where the architecture of the statistical software or the GIS supports only one-to-one (one map and one statistical graph) or one-to-many (one map to multiple statistical graphs) links.

4 Functionality

The main functionality of *DynESDA2* replicates that of its predecessor (see Anselin and Smirnov 1999a, Anselin 2000). It contains maps, histograms, box plots, scatterplots and Moran scatterplots (with the associated computation and permutation-based significance test for Moran’s I statistic) in a framework of dynamically linked windows. As such, both linking and brushing of these graphs is supported. This basic functionality is extended in five main ways.

First, the linking of views has been augmented with the ability to *link multiple maps*, to each other as well as to the statistical graphs. This was accomplished by using MapObjects components for the map rendering and the bit string to tie all windows together. Previously, when working within ArcView, an Avenue script was launched to establish a conversation with the dynamically linked libraries. However, in ArcView’s architecture, this is tied to a single “active” View (the map in ArcView terminology). In order to establish links to a different View, the running script had to first be shut down and then restarted from

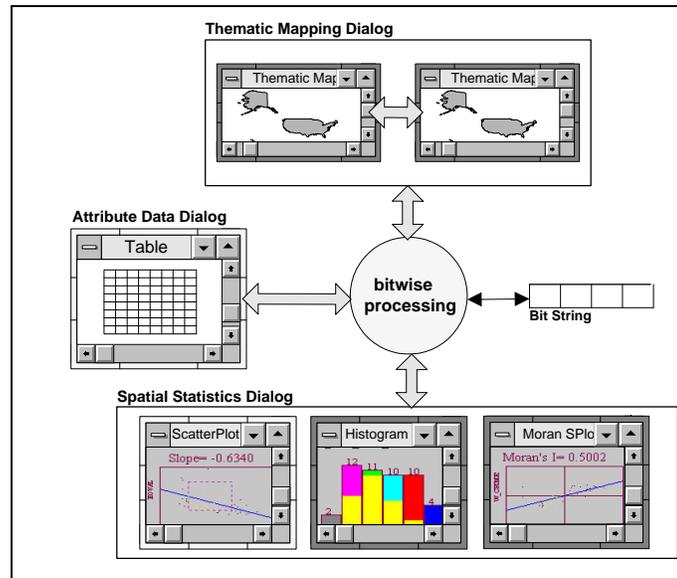


Figure 5: User interaction with multiple linked windows.

the new View. In this process, any links to the original view were lost.⁶ In the current implementation, there is no such constraint.

In addition to linking more than one map with the statistical graphs, multiple maps can now be linked to each other as well. In other words, whenever a feature is highlighted in one of the windows, the corresponding object is highlighted in all of them, irrespective of their nature (maps, graphs or table). Since ArcView 3.x does not support links between different Views, this aspect of dynamic linking cannot be implemented in interfaces built on ArcView.⁷ The use of MapObjects components circumvents this constraint, since the bit string keeps track of the relevant objects. Moreover, since all the maps are tied to the same data table and boundary file, there is no possibility of confusion between different “geographies.”

Second, true *map brushing* is now included. Previously, features could only be selected on a map by means of the built-in ArcView select tools. As a result, in the previous *DynESDA*, true brushing was only possible between statistical graphs and a map, in the sense that the “brush” could be applied to the graph, but not to the map.⁸ By keeping the centroids of the map polygons in memory, and applying the brush to those (and thus indirectly to the polygons themselves), true map brushing capability is obtained. Currently, the brush is implemented as a rectangle shape, but eventually it will be allowed to take on any of the shapes available as select tools (see Figure 4).

Third, a *table* has been added as a new view on the data (in the terminology of dynamically linked windows). This implements some simple data base queries using SQL, such as the selection of specific records and/or specific fields. The rows in the table are linked to the other statistical graphs. There is no brushing per se in the table, but when other graphs

⁶This was not unique to the first version of *DynESDA*, but other implementations of linkages between statistical software and ArcView similarly suffer from this limitation.

⁷ArcView implements a form of dynamic linking between the active View and the corresponding table, as well as a corresponding graph, but not between Views.

⁸True map brushing of polygons (as opposed to points) was originally suggested in Monmonier (1989), but to our knowledge, *DynESDA2* is the first actual implementation of this idea in a software tool.

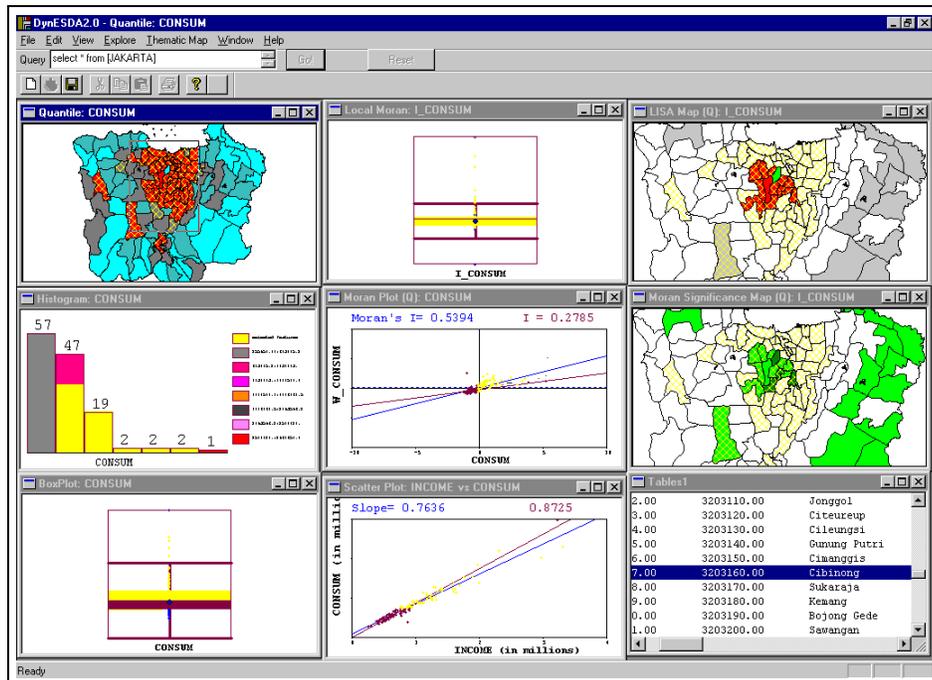


Figure 6: Functionality of DynESDA.

are brushed, the matching records in the table are highlighted.

Fourth, new visualization devices have been included, such as *box maps* and *percentile maps*, as well as a limited degree of animation in the form of a *map movie*. As mentioned before, the mapping functions, which previously relied on the cartographic capability of the ArcView “View” are now implemented using MapObjects components. In addition to standard choropleth maps, such as quantile maps and standard deviational maps (with the facility to zoom in and out), specialized maps have been added that highlight outliers in the data (see Figure 3). Two such devices are the box map, a quartile map with outliers identified, and a percentile map.⁹ Previously, it was necessary to link back and forth to SpaceStat to compute the information needed for these maps. In addition, a type of motion graphic is implemented in the form of a map movie. The map movie is equivalent to an automatic brushing of a box plot from low to high values. Each observation on the map is highlighted in turn, in the order of its magnitude for a selected variable. The map movie can highlight one value at a time, or be cumulative, slowly filling up the map as new values are added. A map movie is a useful device to suggest patterns of spatial heterogeneity in the data (e.g., when all low values are in one region, and higher values in another)

Finally, the visualization of spatial autocorrelation extends the earlier use of the Moran Scatterplot and the computation of Moran's I with a “local” variant, for which the significant locations are shown in a *LISA map* and *Moran Significance Map*. The Local Moran statistic (Anselin 1995) indicates, for each location, the extent to which the pattern of values in neighboring locations is compatible with spatial randomness. Rejection of this null hypothesis suggests either clustering (of high or low values), or spatial outliers (low values surrounded by high, or vice versa). The LISA statistics are visualized in four windows. One shows a map with the locations with significant values for the local Moran, with different

⁹See Anselin (1999a) for a more extensive description of these ESDA tools.

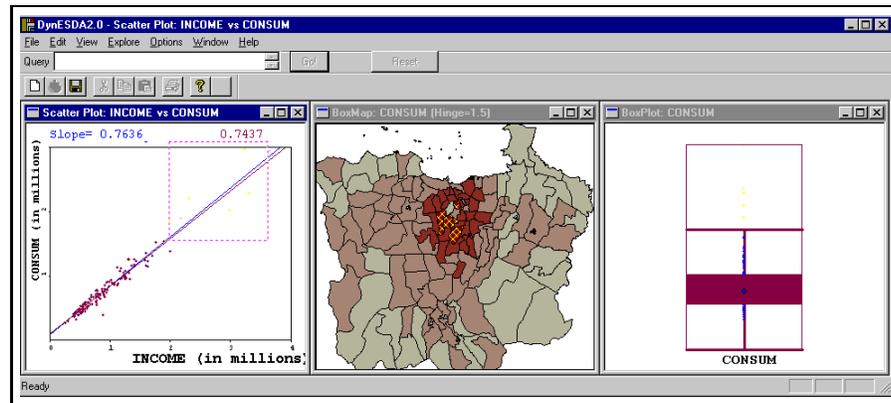


Figure 7: Distribution of household consumption in the Jabotabek region.

colors for different degrees of significance ($p < 0.01$ and $p < 0.001$). This is referred to as a Moran Significance Map (see Figure 2). The second map distinguishes between the four types of local association, but only for the locations with significant LISA statistics. This is referred to as a LISA map. Note that the four types of association also match the four quadrants in the Moran Scatterplot.¹⁰ A third graphic associated with the computation of the LISA consists of a box plot for the individual statistics. On average, the Local Moran equal the global Moran statistic, and the box plot can be used as a diagnostic to assess the extent to which this average is representative of the overall pattern. The final graphic consists of the matching Moran scatterplot. The four windows are linked to all other views, allowing for interactive brushing among different representations of local and global spatial autocorrelation. This comes close to the idea of a “spatial association visualizer” outlined in Anselin (1998).

The full functionality of the *DynESDA2* tool is outlined in Figure 6.

5 Illustration

We briefly demonstrate the design and functionality of *DynESDA2* by means of a simple example, using data on the spatial distribution of household consumption in the greater Jakarta metropolitan area of Indonesia.¹¹ This region is referred to as the Jabotabek metropolitan area, and it comprises 130 subdistricts, covering, in addition to Jakarta, the urban areas of Bogor, Tangerang and Bekasi. This metropolitan region is characterized by internal restructuring, both physical as well as in terms of socio-economic characteristics. The primary pattern is one of decentralization of manufacturing (away from the city center) and concentration of finance and services, changing the nature of urban areas in the core.¹² We consider the spatial distribution of median household consumption in the region, as well as its relation to disposable (household) income in a simple linear consumption function.

In Figure 7, the central graph shows a quartile map for household consumption, with the “outliers” highlighted in the center. The districts in question happen to all be near the core of Jakarta, suggesting the existence of a “cluster.” While these are the locations that match the outlying data points on the traditional box plot on the right, the latter does not

¹⁰For a more technical description, see Anselin (1995, 1996).

¹¹The data are from the 1999 Indonesia socio-economic survey, or SUSENAS (BPS 1999).

¹²For an extensive discussion of the substantive issues involved, see, e.g., Firman (1998).

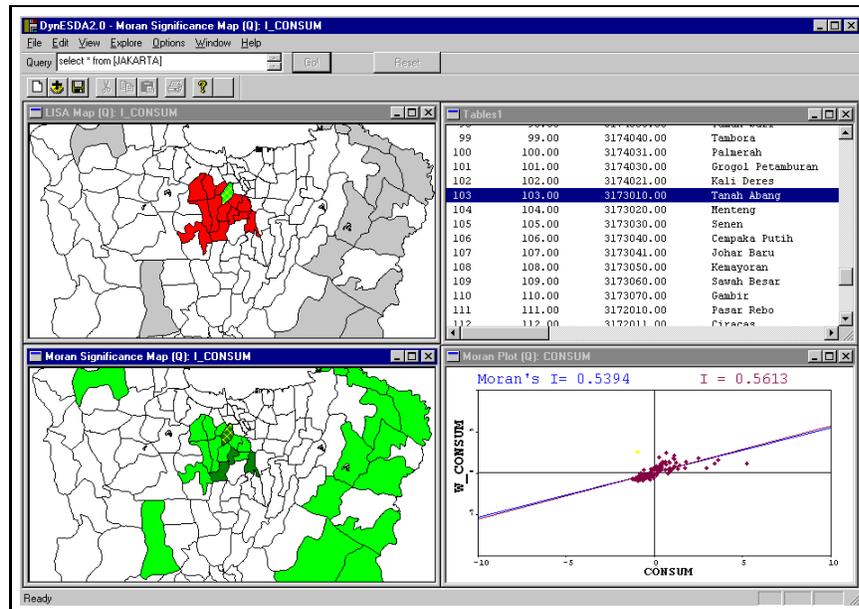


Figure 8: Visualizing local and global spatial autocorrelation.

shed light on any “spatial” pattern as in the box map. On the left, a scatterplot with a linear smoother illustrates the slope of the consumption function and how this is affected when the outlying observations are removed from the analysis (the slope of 0.74 on the upper right hand side of the graph is the value obtained without the outliers). Clearly, the selection in one of the graphs is linked to the matching observations on the other graphs.

We further investigate the spatial pattern of consumption expenditures by focusing on the districts with a significant Local Moran and the type of spatial autocorrelation suggested, as illustrated in Figure 8. The graph on the lower right presents the Moran scatterplot and suggest a (significant) degree of positive spatial autocorrelation (Moran’s I is 0.54). The local autocorrelation provides a more fine grained view of the association however. The significant districts are shown in the lower left map, with the classification by type of association given in the LISA map at the upper left. Both maps have been zoomed in to provide a better view of the pattern near the core city. While the evidence overwhelmingly suggests a significant “cluster” of high household consumption near the center-city, there is one district that does not fit the mold. In the very center, the Tanah Abang district (identified in the Table on the upper right) shows a significant pattern of negative association, i.e., a low value surrounded by high values. A look at the matching point in the Moran scatterplot confirms that this is indeed a location with below average consumption, surrounded by above average neighbors. This pattern is such that it is highly unlikely to occur under spatial randomness. The district in question is indeed a remaining area of urban decay, where the conversion to business activity, characteristic of the other core districts, has not (yet) taken place.

As a final illustration, we consider evidence of spatial heterogeneity, or the significant difference between moments of a distribution and model parameters by regional “regimes.” In Figure 9, a brush is centered on the core urban area, selecting those districts in the map and in the matching box plot. This also *eliminates* the selected points from the computation of Moran’s I and the slope of the consumption function. The latter, shown in the bottom

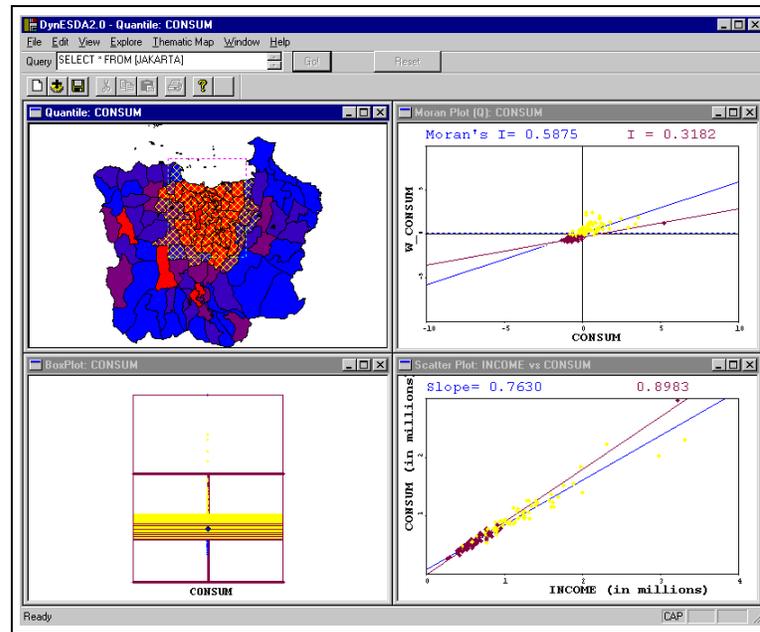


Figure 9: Visualizing spatial heterogeneity.

right graph, is substantially higher than for the data set as a whole (0.90 vs. 0.76), confirming the well-known phenomenon that the propensity to consume (the slope of the regression line) is not constant, but inversely related to income (lower income groups, outside the core districts, have higher propensity to consume). Moving the brush over the map would allow the analyst to interactively assess the degree of change in this parameter over subregions of the data, suggesting candidate “spatial regimes” for further analysis by means of spatial econometric software.

A “static” description such as the one given here rarely gives proper credit to the type of additional insight generated through a dynamic interaction with the data, especially when complemented with substantive knowledge of the issue at hand. Such interaction often provides significant added value over a more traditional data analysis, motivating further work on extending and refining the current framework.¹³

6 Future Directions

DynESDA2 is a work in progress and part of a more comprehensive strategy to facilitate and promote the use of spatial analytical tools in the social sciences (Goodchild *et al.* 2000). Development is ongoing, and several minor refinements to the framework (legends, color choice, print facilities) are in the process of being added.

A major medium-term effort consists of extending the functionality to a broader range of spatial statistics and data exploration tools, as well as to other data structures. The current tools apply to polygons, a special case of lattice (or regional) data. Ongoing development deals with adding points (as lattice data) as well as flows. Information on spatial arrangement for these data structures is no longer based on simple contiguity (constructed from

¹³For a recent example of an instance where the use of ESDA through dynamically linked windows generated “surprising” new insights, see Messner *et al.* (1999).

the boundary files for polygons), but requires the efficient computation and management of distance measures. This is in the process of being pursued.

To date, the goal of full modularity in the form of Microsoft compliant COM “components” has only been partially achieved. Future work will focus on completing the componentization so that the functionality can be leveraged by any com-compliant software, including various GIS and statistical software packages. The end result is envisaged as an open and modular library of components for exploratory spatial data analysis.

References

- Anselin L. 1992. *SpaceStat, a Software Program for Analysis of Spatial Data*, National Center for Geographic Information and Analysis (NCGIA), University of California, Santa Barbara, CA.
- Anselin L. 1995. Local indicators of spatial association — LISA, *Geographical Analysis*, 27: 93–115.
- Anselin L. 1996. The Moran scatterplot as an ESDA tool to assess local instability in spatial association, in Fischer M., Scholten H. and Unwin D. (eds.) *Spatial Analytical Perspectives on GIS in Environmental and Socio-Economic Sciences*, Taylor and Francis, London, 111–125.
- Anselin L. 1998. Exploratory spatial data analysis in a geocomputational environment, in Longley P.A., Brooks S., Macmillan B. and McDonnell R. (eds.) *Geocomputation: A Primer*, John Wiley, New York, NY, 77–94.
- Anselin L. 1999a. Interactive techniques and exploratory spatial data analysis, in Longley P.A., Goodchild M.F., Maguire D.J. and Rhind D.W. (eds.) *Geographical Information Systems: Principles, Techniques, Management and Applications*, John Wiley, New York, NY, 251–264.
- Anselin L. 1999b. *SpaceStat Software Program for Spatial Data Analysis, Version 1.90*, BioMedware, Ann Arbor, MI.
- Anselin L. 2000. Computing environments for spatial data analysis, *Journal of Geographical Systems*, 2: 201–220.
- Anselin L. and Bao S. 1996. *SpaceStat.apr User's Guide*, Working Paper 9628, Regional Research Institute, West Virginia University, Morgantown, WV.
- Anselin L. and Bao S. 1997. Exploratory spatial data analysis: Linking SpaceStat and ArcView, in Fischer M.M. and Getis A. (eds.) *Recent Developments in Spatial Analysis*, Springer-Verlag, Berlin, 35–59.
- Anselin L. and Getis A. 1992. Spatial statistical analysis and geographic information systems, *The Annals of Regional Science*, 26: 19–33.
- Anselin L. and Smirnov O. 1999a. *The DynESDA Extension for ArcView 3.0*, Bruton Center, University of Texas at Dallas, Richardson, TX.
- Anselin L. and Smirnov O. 1999b. *The SpaceStat Extension for ArcView 3.0*, Bruton Center, University of Texas at Dallas, Richardson, TX.
- Anselin L., Dodson R. and Hudak S. 1993. Linking GIS and spatial data analysis in practice, *Geographical Systems*, 1: 3–23.

- Bao S. and Anselin L. 1998. Linking spatial statistics with GIS: Operational issues in the SpaceStat-ArcView link and the S+Grassland link, in *ASA Proceedings of the Section on Statistical Graphics*, American Statistical Association, Alexandria, VA, 61–66.
- Bao S., Anselin L., Martin D. and Stralberg D. 2000. Seamless integration of spatial statistics and GIS: The S-Plus for ArcView and the S+Grassland Links, *Journal of Geographical Systems*, 2: 287–306.
- BPS 1999. *The National-Social and Economic Survey (SUSENAS). Household Survey*, Central Bureau of Statistics, Republic of Indonesia, Jakarta, Indonesia.
- Brundson C. 1998. Exploratory spatial data analysis and local indicators of spatial association with xlistat, *The Statistician*, 47: 471–484.
- Cook D., Majure J., Symanzik J. and Cressie N. 1996. Dynamic graphics in a GIS: A platform for analyzing and exploring multivariate spatial data, *Computational Statistics*, 11: 467–480.
- Cook D., Symanzik J., Majure J.J. and Cressie N. 1997. Dynamic graphics in a GIS: More examples using linked software, *Computers and Geosciences*, 23: 371–385.
- Dykes J.A. 1997. Exploring spatial data representation with dynamic graphics, *Computers and Geosciences*, 23: 345–370.
- Dykes J.A. 1998. Cartographic visualization: Exploratory spatial data analysis with local indicators of spatial association using Tcl/Tk and cdv, *The Statistician*, 47: 485–497.
- Firman T. 1998. The restructuring of the Jakarta metropolitan area: A global city in Asia, *Cities*, 15: 229–243.
- Goodchild M., Anselin L., Appelbaum R. and Harthorn B. 2000. Toward spatially integrated social science, *International Regional Science Review*, 23: 139–159.
- Goodchild M.F., Haining R.P., Wise S. and others 1992. Integrating GIS and spatial analysis — problems and possibilities, *International Journal of Geographical Information Systems*, 6: 407–423.
- Haining R.F., Ma J. and Wise S. 1996. Design of a software system for interactive spatial statistical analysis linked to a GIS, *Computational Statistics*, 11: 449–466.
- Haining R.F., Wise S. and Ma J. 1998. Exploratory spatial data analysis in a geographic information system, *The Statistician*, 47: 457–469.
- Haining R.F., Wise S. and Ma J. 2000. Designing and implementing software for spatial statistical analysis in a GIS environment, *Journal of Geographical Systems*, 2: 257–286.
- Haslett J., Wills G. and Unwin A. 1990. SPIDER — an interactive statistical tool for the analysis of spatially distributed data, *International Journal of Geographic Information Systems*, 4: 285–296.
- Haslett J., Bradley R., Craig P., Unwin A. and Wills G. 1991. Dynamic graphics for exploring spatial data with applications to locating global and local anomalies, *The American Statistician*, 45: 234–242.
- Messner S., Anselin L., Baller R., Hawkins D., Deane G. and Tolnay S. 1999. The spatial patterning of county homicide rates: An application of exploratory spatial data analysis, *Journal of Quantitative Criminology*, 15: 423–450.

- Monmonier M. 1989. Geographic brushing: Enhancing exploratory analysis of the scatterplot matrix, *Geographical Analysis*, 21: 81–4.
- Symanzik J., Majure J. and Cook D. 1994a. Dynamic graphics in a GIS: A bidirectional link between ArcView 2.0 and XGobi, *Computing Science and Statistics*, 27: 299–303.
- Symanzik J., Majure J., Cook D. and Cressie N. 1994b. Dynamic graphics in a GIS: A link between ARC/INFO and XGobi, *Computing Science and Statistics*, 26: 431–435.
- Symanzik J., Megretskaia I., Majure J. and Cook D. 1997. Implementation issues of variogram cloud plots and spatially lagged scatterplots in the linked ArcView 2.1 and XGobi environment, *Computing Science and Statistics*, 28: 369–374.
- Symanzik J., Kotter T., Schmelzer S., Klinke S., Cook D. and Swayne D. 1998. Spatial data analysis in the dynamically linked ArcView/XGobi/XploRe environment, *Computing Science and Statistics*, 29: 561–569.
- Symanzik J., Cook D., Lewin-Koh N., Majure J.J. and Megretskaia I. 2000. Linking ArcView and XGobi: Insight behind the front end, *Journal of Computational and Graphical Statistics*, 9: 470–490.
- Unwin A. 1996. Exploratory spatial analysis and local statistics, *Computational Statistics*, 11: 387–400.
- Wall P. and Devine O. 2000. Interactive analysis of the spatial distribution of disease using a geographic information system, *Journal of Geographical Systems*, 2: 243–256.
- Wilhelm A. and Steck R. 1998. Exploring spatial data by using interactive graphics and local statistics, *The Statistician*, 47: 423–430.
- Wise S., Haining R. and Ma J. 2001. Providing spatial statistical data analysis functionality for the GIS user: the SAGE project, *International Journal of Geographic Information Science*, 15: 239–254.
- Zhang Z. and Griffith D. 1997. Developing user-friendly spatial statistical analysis modules for GIS: An example using ArcView, *Computers, Environment and Urban Systems*, 21: 5–29.
- Zhang Z. and Griffith D. 2000. Integrating GIS components and spatial statistical analysis in DBMSs, *International Journal of Geographical Information Science*, 14: 543–566.