

Visual Analytics for Profiling Land Use Changes

Claudio Santos*, Maryam Hosseini[†], João Rulff[‡], Fabio Miranda[§]

Luc Wilson[¶], Claudio Silva[‡], Nivan Ferreira^{||}, Marcos Lage*

*Universidade Federal Fluminense, [†]Massachusetts Institute of Technology, [‡]New York University

[§]University of Illinois Chicago, [¶]Kohn Pedersen Fox Associates PC, ^{||}Universidade Federal de Pernambuco

claudioj@id.uff.br, maryamh@mit.edu, jlrulff@nyu.edu, fabiom@uic.edu

csilva@nyu.edu, nivan@cin.ufpe.br, mlage@ic.uff.br

Abstract—The growth of cities calls for regulations and zoning rules on how each piece of urban space will be used. Tracking land use can reveal a wealth of information about urban development. For that matter, cities have been releasing data sets describing the historical evolution of the shape and the attributes of land units. The complex nature of land-use data, however, makes the analysis of such data challenging and time-consuming. To address these challenges, we propose URBAN CHRONICLES, a visual analytics system that enables interactive exploration of land-use changes. Using New York City’s *Primary Land Use Tax Lot Output (PLUTO)*, we show the system’s capabilities to explore the data from several years at different scales. URBAN CHRONICLES supports on-the-fly aggregation and filtering operations that leverage the hierarchical nature of the data set to index the shape and attributes of geographical regions that change over time. Finally, we demonstrate the system’s utility through case studies that analyze the impact of Hurricane Sandy on land use attributes and the effects of rezoning plans in Brooklyn.

I. INTRODUCTION

Cities have long been the center of innovation and development. As people concentrate on these hubs of culture and services, hamlets give way to towns and towns to metropolises. This growth calls for regulations and plans to best divide and utilize a precious resource: land. Cities have been adopting and modifying regulations that govern land use and development. Zoning resolutions were defined to facilitate the management of urban spaces, controlling the shape and size of the buildings and public spaces, preserving urban fabrics, and determining how each piece of land can be used [1]. The zoning rules can change to comply with the new demands of urban life. Tracking the zoning codes and land use changes can help draw a precise picture of how urban economic, social, and public policies have evolved and inform stakeholders of the delayed impacts of policy changes and major events such as natural disasters or economic crises.

The complex nature of zoning codes and land use regulations makes analyzing such data challenging and time-consuming. For instance, since such extensive data spans different city agencies, each having its data collection and recording methods, the unified data set is highly prone to errors and inconsistencies. On top of that, land use and zoning data are often large and complex. In New York City (NYC), the data set describing primary land use has over 83 attributes and more

than 320,000 geographical units for the boroughs of Manhattan and Brooklyn. This data can be considered a spatiotemporal data set, but the spatial attribute is not merely composed of points but complex geometric primitives. New geographical units are created, destroyed, split, and merged over time, resulting in attributes and geometry changes. Tracking these changes is, therefore, essential. Most of the existing land use and zoning visualization tools are either merely exploratory tools offering very limited or no analysis capabilities [2] or do not take the temporal aspect of the data into account [3].

To address the challenges involved in the analysis of land use data, we introduce URBAN CHRONICLES, a visual analytics system for profiling city land use data. URBAN CHRONICLES is developed through ongoing collaborations with urban planners and architects working with land use and zoning data as part of their projects. The system implements an in-memory tree-based data structure designed to handle land units’ geometries; and a visual interface that enables the visualization of land use data at different scales giving the user analytical capabilities that would not be possible in a single scale. Our system uses the *Primary Land Use Tax Lot Output (PLUTO)* data, a spatiotemporal data set that describes the land use information in NYC since 2002. We use it as an example to introduce URBAN CHRONICLES, although it could also be used to explore other land use inventories. We demonstrate the effectiveness of the system through case studies set in NYC. They analyze the impact of the National Flood Insurance Program of 2013, a year after Hurricane Sandy, on residential development activities in lower-income neighborhoods and evaluate the Downtown Brooklyn rezoning plan to absorb the office space demand generated by Financial District. To summarize, our contributions are:

- We introduce URBAN CHRONICLES, a web-based system that enables urban planners and architects to visually explore large data sets composed of geographical, political, and administrative land unit information.
- We build a general-purpose tree-based data structure to leverage the hierarchical nature of land use data. The data structure enables interactive time when querying large spatiotemporal land use data sets.
- We highlight the utility of URBAN CHRONICLES through two case studies in NYC and performed by urban planners and architects.




Fig. 1. Lot changes in PLUTO: The geometry of the lot was updated to represent its real geometry (2014) better; In the following year, it was split, creating a new lot (2015); In 2016, the larger lot became a mixed-use area; and in 2017 the two lots were merged.

II. RELATED WORK

Land use and zoning analysis. Scientists in different fields have studied land use and land cover change. Environmental scientists looked at how such modifications can impact climate change [4], while urban planners showed that it could signal gentrification [5], [6] and affect different urban indices [7]–[9]. Recently, data sets, such as PLUTO [10] and Chicago’s Land Use Inventory [11], were used in several studies. Still, most consider just a few attributes, short periods, and coarse spatial granularity [12]–[14].

Interactive visualization. To be effective, visual analytics systems should strive for *interactivity*, with response times below 500ms [15]. Thus, specialized data structures [16]–[18] and parallelism techniques were explored over the past years [19], [20]. These works do not consider data containing geometric shapes that evolve. Therefore, they are not suitable for data representing land use data.

Urban visual analytics. In the past decade, various tools and systems have been proposed to analyze and explore urban data [21]–[24], helping experts to gain a deeper understanding of cities, evaluate policies, and plan developments [25]. Several tools have also been proposed to handle complex urban geometry, such as buildings [26], [27]. Although their goal is to inform the decision-making process, these systems were not developed as general-purpose tools to cover the whole spectrum of geographical-based analysis. For more focused tasks, such as land use analysis, these tools do not provide the capability to interactively explore and track the temporal variation of the data sets.

III. BACKGROUND

Lots & rezoning. There are multiple standards to locate and identify each piece of land. The system used in the U.S. and Canada is called *Lot and Block Survey System*, under which blocks are defined as “A tract of land bounded on all sides by streets, or a combination of streets, public parks, railroad rights-of-way, pierhead lines or airport boundaries”. Blocks are then subdivided into lots, the smallest parcel of land that can be owned, purchased, and sold. In NYC (the running example used in this paper), lots are identified with a unique code resulting from the concatenation of borough, block, and lot numbers corresponding to their location, called the BBL code. A lot (see Figure 1) can be divided to create new lots, or multiple lots can merge to create a larger lot. The zoning

designation can change through a process called *rezoning*. It often seeks to either increase the density and encourage development by setting *less restrictive* rules on the height, size, design, and use of the buildings; or to prevent density increase, protect local businesses and the historic fabric of the neighborhood by setting *more restrictive* rules.

The PLUTO data. PLUTO is an invaluable and extensive land use and geographic data set, describing almost every piece of land in NYC [10]. MapPLUTO is derived from merging the PLUTO attributes with the Tax Lot Polygon features from the Department of Finance. Among others thing, the data set provides information regarding assessed land value and primary building features. PLUTO and MapPLUTO are updated twice a year. We collected PLUTO data from 2002 to the first semester of 2017. During this period, 22 versions (18GB) of the data set were released. Each PLUTO data release is a collection of shapefiles, one for each of the five boroughs in NYC, containing the geometries and attributes of the lots. The raw historical data contains many redundancies that can be explored for storage optimization. For example, the geometry of most of the lots in NYC remains unchanged for several years. Figure 1 illustrates all changes that the geometry and attributes of a lot can undergo. We stress that tracking and visualizing the changes of all lots of a city is a challenging task that we aim to solve with URBAN CHRONICLES. Moreover, the attributes of the lots have gone through several changes over the years. Examples of issues one observes when trying to consolidate the different PLUTO releases are changes in attribute names, changes in attribute definition, and inaccurate attribute information. PLUTO, NYC’s data set that is our motivating example, is a spatiotemporal data whose spatial component of the elements are geographical regions, where each region is described by its polygonal boundary and a list of attributes that change over time. Next, we describe the main features of PLUTO.

IV. URBAN CHRONICLES SYSTEM

In our ongoing collaboration with domain experts, we identified a set of tasks that they are interested in performing and defined the following requirements for URBAN CHRONICLES: **[R1] Enable data exploration at different city scales.** Enable the exploration of land use data at different geographical levels. An urban planner might use the neighborhood level to identify areas undergoing a process of gentrification and then analyze the phenomena at the block or lot level.

[R2] Enable the exploration of attribute changes over time.

Allow users to explore the rate of change in the attributes over time while enabling the identification of outliers and patterns.

[R3] Enable the exploration of lots of interest. Filtering lots based on tax, land use, or zoning attributes allows experts to identify areas that follow certain criteria.

[R4] Enable the exploration of temporal changes of lots.

Keeping track of changes in the number of lots in a neighborhood is crucial. The increase in lot mergers can be a signal of an incentive zoning program or transferable development rights, which can have a large impact on the community.

[R5] Support interactive query times. All the previous tasks should be executed interactively since a response time greater than 500ms can significantly impact visual analysis [15].

To meet these requirements, URBAN CHRONICLES has three modules: the *Data Storage* (Section IV-A), the *Query Processor* (Section IV-B) and the *Visual Interface* (Section IV-C).

A. Data Storage

Public administration uses several different geographical subdivisions to define city policies, rezoning plans, etc. For this reason, URBAN CHRONICLES provides users the ability to explore PLUTO data based on the borough, neighborhood, community district, block, and lot levels (requisite **R1**). This requisite demands the construction of a spatial index that stores the lots belonging to each level’s geographical regions. To do so, we built a *Data Storage* module composed of in-memory *data containers* and a *spatiotemporal index*. The storage works together with the *Query Processor* module to handle requests from the *Visual Interface* module (see Section IV-C)

The *data containers* use lists to store the lots’ and geographical regions’ geometries, as well as their attributes. The *Spatiotemporal Index* is implemented using a tree containing five levels: the city, boroughs, neighborhoods (or community districts), blocks, and lots. We observe that either the neighborhoods or the community district regions are considered during the data structure construction. Since this option has no impact on the structure’s design, we will assume the neighborhood regions are chosen from now on. Also, note that a geographical region in level $l + 1$ is always contained in a single region of level l . In this way, the tree’s root node represents the *city* level and covers the entire NYC. Analogously, the *borough* level contains nodes that represent each of the city’s five boroughs; the *neighborhood* level includes nodes that represent the neighborhoods of each borough; and so on. Each leaf node stores one reference to its *Geometries Container* and *Attributes Container*, storing the lot’s data of all PLUTO releases.

B. Query Processor

The query processor uses the *Data Storage*’s spatiotemporal index and containers to process the queries the visual interface produces. The queries supported by the system can be classified into geometry retrieval, attribute aggregation, and lot filtering operations. The geometry retrieval queries return the shape of a set of geographic regions in a given year. For example, it allows retrieving the geometry of the lots that

belong to the block with id 01036 in the Brooklyn neighborhood called Park Slope in 2006. The attribute aggregation queries return an attribute of interest aggregated over a set of geographic regions. For example, it enables the computation of the sum of the area of the lots in each block of the Manhattan neighborhood called SoHo in 2009. The lot filtering operations allow selecting lots based on their attributes. For example, it enables identifying residential or commercial lots. Using a notebook equipped with a Core i7@2.2 GHz, 16 GB of RAM, SSD drive, and NVidia GeForce GTX 1060 GPU, the system was able to compute aggregation and filter queries for the entire Manhattan in 300ms on average (requisite **R5**).

C. Visual Interface

URBAN CHRONICLES’s visual interface allows users to interactively explore the history of PLUTO data at different spatial scales. Figure 2 shows its components, classified as auxiliary and analytical based on their functionalities.

Auxiliary components. The *Main Menu*, the *Configuration Panel*, and the *Color Maps Legend* are labeled as (a), (b), and (c) in Figure 2. The *Main Menu* allows the user to show/hide the other components. The *Configuration Panel* lets the user adjust the system options, *i.e.*, it is possible to define what borough the user wants to explore; what set of geographical regions (neighborhoods or community districts) should be used; whether the block levels may be skipped; and the color mapping used in the *Map View* (see details next).

Analytical components. The *Data History View*, the *Map View*, and the *Lot Filtering* components are labeled as (d), (e), and (f) in Figure 2. The *Data History View* contains two linked visualizations, a line chart, and a heat matrix. They enable the user to explore attributes of a lot of interest and their change over the years (requirement **R2**). The user can select the attribute visualized in both charts using the top toolbar. Each line of the heat matrix represents a geographical region of the selected level. For example, if the neighborhood level is active, each line of the matrix represents one neighborhood of NYC. Also, each matrix column is associated with the release year of the PLUTO data. By default, the lines of the matrix are sorted using the alphabetical order of the names of the regions, but they can also be sorted based on the values of a given year. In the example of Figure 2, the heat matrix is sorted by the values of 2009.2. Using the top toolbar, the color of each square can be set to encode the variation of the selected attribute over a fixed year interval. Whenever the value of a square is invalid, or the lot doesn’t exist in the associated year, it is colored gray. The heat matrix allows users to simultaneously analyze an attribute’s variation (or its value) in all regions/years and facilitates the identification of prominent values. In Figure 3, we observe a significant drop in land values of several Brooklyn neighborhoods, blocks, and lots in 2013, a year after Hurricane Sandy. Similarly, each line in the line chart represents a region of the current exploration level. In Figure 2, the line chart contains two lines representing Lincoln Square (blue) and Hudson Yards (orange) neighborhoods. The chart’s horizontal axis spans the



Fig. 2. The components of the URBAN CHRONICLES' *Visual Interface*.

release years of the PLUTO data. The vertical axis can be configured to represent the attribute values or their yearly variation. The line chart and the heat matrix always display complementary information. If the heat matrix is configured to show the attribute variation, the line chart is automatically set to show the attribute values and vice versa. The line chart allows users to compare values of regions of interest over the years easily. The line charts in Figure 3 show the variation in the land value for selected regions in the neighborhood, block, and lot levels. The value drop can be seen on all scales. The *Map View* shows the set of geographical regions of the current exploration level in one of the years (requirements **R3** and **R4**). In Figure 2, the *Map View* shows the Manhattan neighborhoods colored using the 2009.2 data. Regions can be colored based on attribute values or their variation. Figure 2 shows regions colored by the attribute variation. Visualizing geographical regions using the map allows the user to identify spatial patterns that cannot be observed using the heat matrix or the line chart. For instance, in Figure 3, if we inspect the *Map View*, we discover that the drop in the land values happens in the shoreline neighborhoods. On the block and lot levels, it becomes evident that the change happened to blocks and lots located in the peripheral areas closer to the shoreline. The last analytical component is the *Lot Filtering* panel. It allows users to define filters based on attributes. The user can define multiple filters and combine them using boolean operations. As shown in Figure 2, a bar chart with the distribution of the selected attribute values must be brushed to define a filter.

Geographical level navigation. By default, the *Visual Interface* starts in the neighborhood (or community district level) exploration level. However, the user can drill down to the block and lot level to explore details in a specific geographical region. To do so, the user must click on the label of one of the geographical regions in the vertical axis of the heat matrix. Once a label is selected, the *Data History View* and the *Map View* are updated to display the data from the selected region. In Figure 3, the user started the analysis by exploring Brooklyn neighborhoods. After the land value drop was detected, the user clicked on the label of the Canarsie neighborhood. During the inspection of the block level, it became clearer that the drop

was concentrated in the shoreline blocks. By clicking on the label of the blocks, the user could observe the most affected lots. Whenever the user starts to navigate through the different scales, the bottom toolbar of the *Data History View* displays the analysis path that the user is following (Figure 2(d)). This way, once the block or lot level is active, the navigation path can roll up to coarser levels.

V. CASE STUDIES

A. Post Sandy Redevelopment and Affordability

Hurricane Sandy hit NYC in 2012, and the impact was devastating. In 2013, the National Flood Insurance Program (NFIP) required homeowners in high-risk zones to pay insurance premiums based on their flood risk. The program extended the geographical location of high-risk areas. While this can help mitigate the future adverse effects, such provisions pressure lower-income communities and working-class homeowners in high-risk flood areas.

Neighborhood level. We first found a noticeable change in land value in Brooklyn through the *Data History View*. A sharp decrease in the second half of 2013, one year after Sandy, can be easily detected (Figure 3(a)). This decrease happened in shoreline neighborhoods impacted the most by the hurricane. Next, we use URBAN CHRONICLES to explore two affected areas with significantly different normalized total values rates while both were located in high-risk flood zones. The normalized total value gives a more accurate measurement of cross-comparison between regions of different sizes. We chose North Side-South Side, a neighborhood with a total value of 1.45 million dollars per sq. ft, and Canarsie, a lower-income neighborhood with 148 thousand dollars in total value per sq. ft. as of 2013. Figure 3(a) shows their location and normalized total values.

Block level. We drilled down to the block level to explore the pattern of value drop in residential areas within both neighborhoods. As shown on the map (Figure 3(b,c)), while in Canarsie, we see a significant decrease in the normalized total value of the majority of the residential blocks (highlighted in shades of red); in North Side-South Side, some blocks still had their values appreciated (highlighted in shades of blue). One

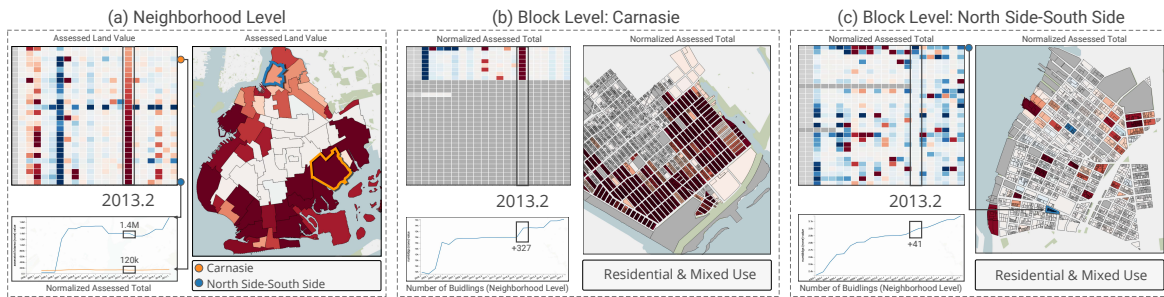


Fig. 3. (a) Neighborhood Level: Temporal evolution of the land value of Brooklyn neighborhoods. (a-bottom left) The line chart shows a sharp difference between the normalized total value of Carnasie (Orange line) and North Side-South Side (Blue line). (b,c - top) Block level: North Side-South Side has some blocks with an increased normalized total value, while in Canarsie, we see a severe value drop in most residential blocks. Line charts depict that 327 residential buildings were constructed in Canarsie (b-bottom left), while North Side-South Side only had 41 new buildings in 2013 (c-bottom left).

potential explanation can be the larger extent of Sandy damage in Canarsie due to lower flood resiliency in residential building construction and lack of coastal protection [28]. This indicates how natural disasters can create more burdens for lower-income neighborhoods. Reports show that only five percent of the residential buildings in Canarsie had flood insurance at the time of Sandy [28]. The new provision would require all homeowners in the flood zone areas to have flood insurance, and the premium would rise by 18% per year. It can be challenging for many households with fixed incomes to afford the increase, which can lead to their displacement. The change in the number of new buildings within the residential areas shows that in 2013, 327 new buildings were constructed in Canarsie Figure 3(b-bottom left), while in North Side-South Side, only 41 new buildings were constructed Figure 3(c-bottom left). A large number of new constructions, specifically the addition of 224 new homes in 2016, can signal a rise in foreclosure and short sales in Canarsie, possibly due to the inability of the majority of uninsured homeowners to accommodate the rising mandatory flood insurance together with the repair costs [29]. Further research can uncover the reasons behind the heightened development to explain the different rebuilding rates.

B. Downtown Brooklyn Rezoning

After 9/11, workers were thought to be unwilling to return to the Financial District (FD) in NYC. So, the city needed to create a new commercial district to keep jobs from moving to other cities. In this scenario, Brooklyn’s plan to transform the Special Downtown Brooklyn District (SDBD) (Figure 4(a)) into a vibrant Business District seemed very plausible [30]. The plan was approved in 2004, and it was predicted that by 2013 Downtown Brooklyn would (1) construct 4.6 million sq. ft. of office space, (2) create 0.9 million sq. ft. of residential space, (3) increase tax revenue, and (4) increase the public space and cultural amenities.

In a report published in 2016, the outcomes of the rezoning plan were analyzed, and it was concluded that it largely had diverged from its initial goals [31]. The anticipated addition of residential space was met way before the planned time, as the line chart shows (Figure 4(c)), and it grew beyond predictions. As stated in the report, by 2016, only 1.3 million

sq. ft of commercial space (including office space) had been developed in the SDBD [31]. Using URBAN CHRONICLES, we could look closer at the office space development trend. The line chart in Figure 4(b) shows that the total office area in Downtown Brooklyn increased in the first four years of the program, but after that, it declined to the point that in 2017, it was less than its initial value in 2004. By comparing the characteristics of the office buildings in FD and Midtown Business District (MBD), which absorbed a significant portion of the demand from FD [32], we could better understand the type of office space in demand. One of the defining characteristics of the FD office space was its large floor plate [33]. As Figure 4(e) shows, in FD, office buildings were constructed on lots that are, on average, 25,000 sq. ft., whereas this number was around 8,000 for SDBD (d-top), more than three times lower. Although MBD also has offices built on smaller lots of 12,500 sq. ft. on average (e-top), it is still significantly higher than SDBD. Moreover, the average office area shows that each lot in MBD has, on average, around 220,000 sq. ft. of area spread across different floors (e-bottom). In SDBD, however, the average office area per lot is around 73,000 sq. ft. (d-bottom) which is significantly smaller in comparison with MBD and FD. Hence, we can say that lot size can be one of the main factors in explaining why SDBD did not successfully absorb the office demand. To accommodate such demand, there is a need for the merging of lots. Besides, the relatively small lot size is a perfect choice for private developers to place their next residential buildings, and the upward shown in Figure 4(c-bottom) confirms this. By examining the map of office and residential buildings in SDBD, Figure 4(b-bottom), we can see that not only the lot sizes are limited, but they also have irregular shapes, making merging possibilities more difficult. The results signify the need for regular reevaluation of the employed strategies, which can open the possibility of making early adjustments to ensure that the initial goals are realistic and attainable and the undertaken strategies lead to the desired outcomes. As we showed, URBAN CHRONICLES can provide a solid basis for such a study.

VI. CONCLUSION AND FUTURE WORK

In this paper, we proposed URBAN CHRONICLES, designed in close collaboration with experts, to enable the interactive vi-

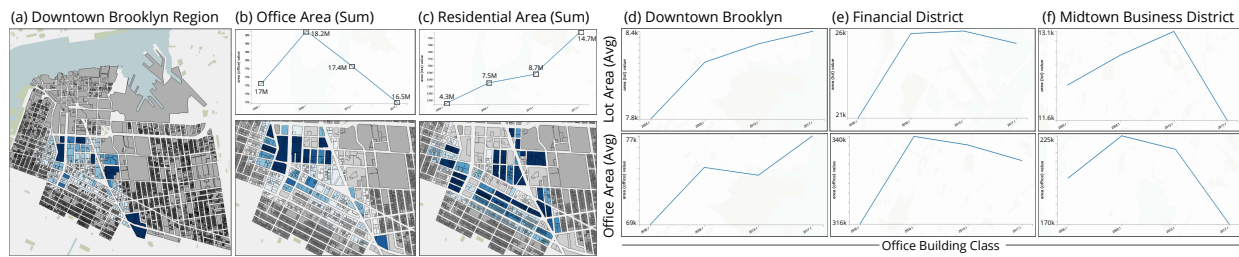


Fig. 4. Using URBAN CHRONICLES to select lots inside the Special Downtown Brooklyn District (SDBD) (a). Top: Total area of offices (b) and residential (c) buildings in SDBD from 2005 to 2017. Bottom: office (b) and residential lots (c) within SDBD, highlighted in white and shades of blue; darker color indicates a larger total area. Line charts depicting the average lot area (top) and office area (bottom) of SDBD, FD, and Midtown Business District (d,e,f).

visual exploration of the versions of the PLUTO data set released from 2002 to 2017. As we show through two case studies, the system allows for advanced analysis of land use patterns and rezoning policies. We intend to extend the system's use to other cities with similar data sets. Although the ideas presented in this paper can be used with any spatiotemporal data set with complex geometries, other challenges can emerge based on the characteristics of the new data. Finally, building a visualization system to simulate and explore rezoning plan scenarios can be an interesting future research topic.

REFERENCES

- [1] E. M. Bassett, *Zoning: The laws, administration, and court decisions during the first twenty years*. Russell Sage Foundation, 1940.
- [2] New York City Department of City Planning (NYC DCP), "ZoLa," Available: <https://zola.planning.nyc.gov/>, 2021.
- [3] The Municipal Art Society of New York, "Accidental Skyline: Air Rights," Available: https://www.mas.org/interactive_features/accidental-skyline-air-rights/, 2021.
- [4] E. Kalnay and M. Cai, "Impact of urbanization and land-use change on climate," *Nature*, vol. 423, no. 6939, pp. 528–531, 2003.
- [5] H. Lim, J. Kim, C. Potter, and W. Bae, "Urban regeneration and gentrification: Land use impacts of the cheonggye stream restoration project on the seoul's central business district," *Habitat International*, vol. 39, pp. 192–200, 2013.
- [6] C. Hamnett and D. Whitelegg, "Loft conversion and gentrification in london: from industrial to postindustrial land use," *Environment and Planning A*, vol. 39, no. 1, pp. 106–124, 2007.
- [7] L. L. M. Goldberg, "Game of zones: neighborhood rezonings and uneven urban growth in bloomberg's new york city," Ph.D. dissertation, Massachusetts Institute of Technology, 2015.
- [8] S. Getz, "Examining the extent to which affordable housing development acts as a catalyst for neighborhood economic development," Ph.D. dissertation, Columbia University, 2017.
- [9] J. Ma, J. C. Cheng, F. Jiang, W. Chen, and J. Zhang, "Analyzing driving factors of land values in urban scale based on big data and non-linear machine learning techniques," *Land Use Policy*, vol. 94, 2020.
- [10] New York City Department of City Planning, "PLUTO README Document," Available: <https://on.nyc.gov/3DOLHYx>, 2021.
- [11] Chicago Metropolitan Agency for Planning (CMAP), "Land Use - CMAP," Available: <https://www.cmap.illinois.gov/data/land-use>, 2021.
- [12] "UrbanReviewer," Available: <http://www.urbanreviewer.org/>, 2021.
- [13] "NYCommons," Available: <https://nycommons.org/>, 2021.
- [14] "Second city zoning," Available: <https://secondcityzoning.org/>, 2021.
- [15] Z. Liu and J. Heer, "The effects of interactive latency on exploratory visual analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2122–2131, 2014.
- [16] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis, "The case for learned index structures," in *Proceedings of the 2018 International Conference on Management of Data*, ser. SIGMOD '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 489–504.
- [17] L. Lins, J. T. Klosowski, and C. Scheidegger, "Nanocubes for real-time exploration of spatiotemporal datasets," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2456–2465, 2013.
- [18] F. Miranda, L. Lins, J. T. Klosowski, and C. T. Silva, "Topkcube: A rank-aware data cube for real-time exploration of spatiotemporal data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 3, pp. 1394–1407, 2018.
- [19] Z. Liu, B. Jiang, and J. Heer, "imMens: Real-time visual querying of big data," *Computer Graphics Forum*, vol. 32, no. 3, pp. 421–430, 2013.
- [20] E. T. Zacharatou, H. Doraiswamy, A. Ailamaki, C. T. Silva, and J. Freire, "GPU rasterization for real-time spatial aggregation over arbitrary polygons," *Proceedings of the VLDB Endowment*, vol. 11, no. 3, pp. 352–365, 2017.
- [21] G. Andrienko, N. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel, "Scalable analysis of movement data for extracting and exploring significant places," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 7, pp. 1078–1094, 2013.
- [22] N. Ferreira, J. Poco, H. T. Vo, J. Freire, and C. T. Silva, "Visual exploration of big spatio-temporal urban data: A study of New York City taxi trips," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2149–2158, 2013.
- [23] N. Andrienko, G. Andrienko, F. Patterson, and H. Stange, "Visual analysis of place connectedness by public transport," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 8, pp. 3196–3208, 2020.
- [24] F. Miranda, H. Doraiswamy, M. Lage, L. Wilson, M. Hsieh, and C. T. Silva, "Shadow accrual maps: Efficient accumulation of city-scale shadows over time," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 3, pp. 1559–1574, 2019.
- [25] Y. Zheng, W. Wu, Y. Chen, H. Qu, and L. M. Ni, "Visual analytics in urban computing: An overview," *IEEE Transactions on Big Data*, vol. 2, no. 3, pp. 276–296, 2016.
- [26] N. Ferreira, M. Lage, H. Doraiswamy, H. Vo, L. Wilson, H. Werner, M. Park, and C. Silva, "Urbane: A 3D framework to support data driven decision making in urban development," in *2015 IEEE conference on visual analytics science and technology (VAST)*. IEEE, 2015, pp. 97–104.
- [27] H. Doraiswamy, E. Tzirita Zacharatou, F. Miranda, M. Lage, A. Ailamaki, C. T. Silva, and J. Freire, "Interactive visual exploration of spatiotemporal urban data sets using Urbane," in *Proceedings of the 2018 International Conference on Management of Data*, ser. SIGMOD '18. New York, NY, USA: ACM, 2018, pp. 1693–1696.
- [28] Department of City Planning, "Resilient neighborhoods canarsie," Available: <https://on.nyc.gov/3Oqopx5>, 2017.
- [29] "Rising tides, rising costs: Flood insurance and New York City's affordability crisis," Available: <https://bit.ly/44Pm4Tf>, 2014.
- [30] New York City Department of City Planning (NYC DCP), "Downtown brooklyn," Available: <https://on.nyc.gov/3KtN0Y>, 2004.
- [31] E. L. Adams, "A decade later in downtown brooklyn: A review of the 2004 rezoning," Available: https://brooklyn-usa.org/wp-content/uploads/2016/02/Downtown-Brooklyn-2004-Rezoning_Final.pdf, 2016.
- [32] A. F. Haughwout, B. Rabin *et al.*, "Exogenous shocks and the dynamics of city growth: Evidence from new york," *Federal Reserve Bank of New York Economic Policy Review*, vol. 11, no. 2, pp. 61–73, 2005.
- [33] F. Fürst, "Empirical analysis of office markets: A spatio-temporal approach," 2006.