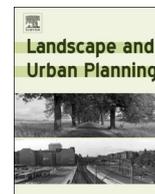




University of  
Massachusetts  
Amherst

## Modeling non-stationary urban growth: The SPRAWL model and the ecological impacts of development

Item Type	article;article
Authors	McGarigal, Kevin;Plunkett, Ethan;Willey, Lisabeth L.;Compton, Brad;DeLuca, Bill;Grand, Joanna
DOI	<a href="https://doi.org/10.1016/j.landurbplan.2018.04.018">https://doi.org/10.1016/j.landurbplan.2018.04.018</a>
Rights	UMass Amherst Open Access Policy
Download date	2024-09-06 19:17:08
Item License	<a href="http://creativecommons.org/licenses/by-nc-nd/4.0/">http://creativecommons.org/licenses/by-nc-nd/4.0/</a>
Link to Item	<a href="https://hdl.handle.net/20.500.14394/37660">https://hdl.handle.net/20.500.14394/37660</a>



## Research Paper

# Modeling non-stationary urban growth: The SPRAWL model and the ecological impacts of development

Kevin McGarigal\*, Ethan B. Plunkett, Lisabeth L. Willey<sup>1</sup>, Bradley W. Compton, William V. DeLuca, Joanna Grand<sup>2</sup>

Department of Environmental Conservation, University of Massachusetts, Amherst, MA 01003, USA

## ARTICLE INFO

## Keywords:

Land use land cover change  
Scenario analysis  
Urban development  
Ecological integrity

## ABSTRACT

Urban development is a principal driver of landscape change affecting the integrity of ecological systems and the capacity of the landscape to support species. We developed an urban growth model (SPRAWL), evaluated it with hindcasting, and used it to simulate urban growth across the northeastern United States between 2010 and 2080 under four alternative scenarios. In the model, urban growth is constrained by demand for new development for each time step at the subregional scale. Demand is subsequently allocated to local application panes (5 km on a side within 15 km window) using a unique landscape context matching algorithm, such that the more historical development that occurred in the matched training windows the higher the proportion of future demand assigned to the pane. Lastly, demand in each pane is allocated among development types and then allocated to individual patches based on suitability surfaces unique to that landscape context. SPRAWL has a multi-level, multi-scale structure that captures urban growth drivers operating at multiple scales and, when combined with the unique matching and suitability algorithms, induces non-stationarity in urban growth across time and space. Our evaluation indicated that SPRAWL was highly discriminatory, well-calibrated, and highly predictive of new development, but performed weakly for redevelopment transitions. We evaluated the ecological impacts of four alternative urban growth scenarios varying in total demand for new development and “sprawliness” of new development relative to historical patterns using an ecological integrity index. The results were consistent with expectations and demonstrated the potential of SPRAWL for scenario analysis.

## 1. Introduction

Urban growth (i.e., land use change associated with residential, commercial or industrial development) is a major landscape change driver in many parts of the world, and in many areas it is the principal driver of permanent landscape change and habitat loss affecting both the integrity of ecological systems and the capacity of the landscape to support biodiversity (Allen 2006; Grimm et al. 2008; Newbold et al. 2015). This is especially the case in the northeastern United States (hereafter simply the Northeast) where the multi-century history of human land use associated with urban development has caused major changes in the capacity of the landscape to support biodiversity (e.g., Foster & Aber 2004) and continues to do so today (DeNormandie & Corcoran 2009). To combat this ongoing loss of biodiversity it is paramount that we better understand and manage urban growth trajectories. To this end, conservation organizations and others are

increasingly seeking ways to predict future urban growth patterns in order to design conservation strategies that proactively protect important natural areas and the fish, wildlife, and other components of biodiversity that they support (e.g., Sim & Mesev 2014; Wear, Pye, & Riitters 2004).

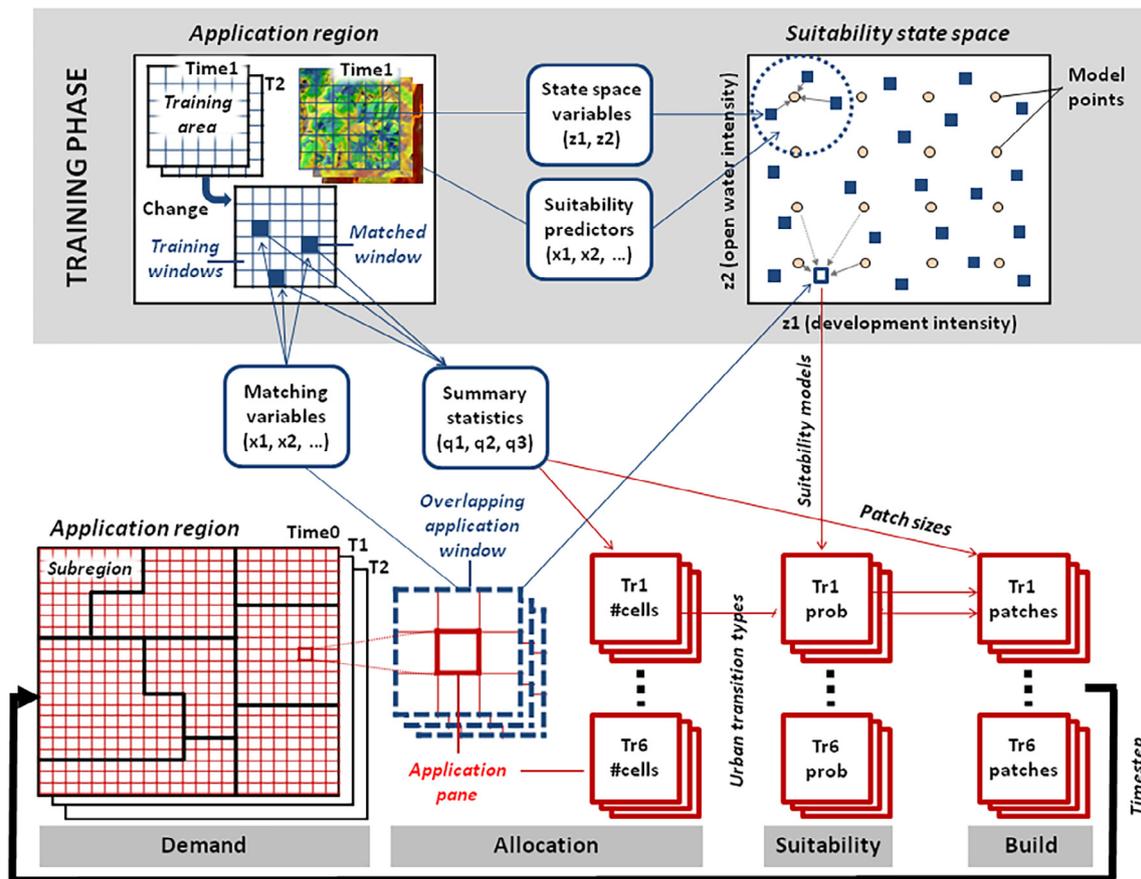
The Designing Sustainable Landscapes (DSL) project (McGarigal, Compton, Plunkett, Deluca, & Grand 2017) was established under the auspices of the North Atlantic Landscape Conservation Cooperative (NALCC) and Northeast Climate Science Center (NECSC) to develop a modeling framework for predicting future landscape changes in the Northeast driven by urban growth and climate change (among other processes), assessing the potential ecological impacts of those landscape changes, and designing landscape conservation strategies to combat those potential impacts. A key component of this modeling framework necessarily addressed urban growth. Specifically, we sought an urban growth model that met the following criteria:

\* Corresponding author.

E-mail addresses: [mcgarigalk@eco.umass.edu](mailto:mcgarigalk@eco.umass.edu) (K. McGarigal), [eplunkett@eco.umass.edu](mailto:eplunkett@eco.umass.edu) (E.B. Plunkett), [lwilley@antioch.edu](mailto:lwilley@antioch.edu) (L.L. Willey), [bcompton@eco.umass.edu](mailto:bcompton@eco.umass.edu) (B.W. Compton), [wdeluca@eco.umass.edu](mailto:wdeluca@eco.umass.edu) (W.V. DeLuca), [jgrand@audubon.org](mailto:jgrand@audubon.org) (J. Grand).

<sup>1</sup> Department of Environmental Studies, Antioch University, Keene, NH 03431-3516, USA.

<sup>2</sup> Science Division, National Audubon Society, Northampton, MA 01060, USA.



**Fig. 1.** Schematic outline of the SPRAWL model. (1) **Training.** Training data (which need not coincide with application data) consists of land cover in two time steps. The training landscape is divided into 15 km windows within each subregion. Training windows are placed in two-dimensional suitability state space based on their intensity of development and open water. Logistic regression models are built for each uniformly distributed model point within the state space to determine suitability. (2) **Demand.** The expected future development rate is supplied by subregion (we used counties or census block statistical areas) from outside the model (we used RPA assessments). (3) **Allocation.** The demand within a subregion is allocated to each 5 km application pane (within an overlapping 15 km application window) based on the relative historical allocation to the three most similar (matching) training windows. (4) **Suitability.** Each application pane (within an overlapping window) is placed in the suitability state space, and nearby logistic regression models are averaged to determine suitability at the cell level. (5) **Building.** Finally, patches of development (for each of 6 different transition types) are built within each application pane based on the corresponding suitability surface, matched to the distribution of patch sizes for matching windows in the training data.

- Ability to efficiently model land cover transitions at the 30 m resolution aggregated into realistic patches over the entire Northeast;
- Empirically-driven to the extent possible to minimize subjective parameterization and avoid having to make assumptions about the socioeconomic mechanisms that drive urban growth (as needed in a process-driven approach);
- Ability to model multiple types of urban land use and land-use changes associated with urban growth;
- Ability to implement scenario-based analysis to reflect our uncertainty in the future urban growth rate and spatial patterns of development; and
- Ability to simulate non-stationarity in the amount and spatial pattern of urban growth across both space and time.

There have been myriad efforts to model land use changes associated with urban growth, and many excellent reviews (e.g., [Berling-Wolff & Wu, 2004a](#); [Li & Gong, 2016](#); [Matthews, Gilbert, Roach, Polhill, & Gotts, 2007](#); [Parker, Manson, Janssen, Hoffmann, & Deadman, 2003](#); [Santé, García, Miranda, & Crecente, 2010](#); [Triantakonstantis & Mountrakis, 2012](#); [Verburg, Schot, Dijst, & Veldkamp, 2004](#)). Despite marked variations in model structure, data requirements and application scales, urban growth models can be unified in a general framework of urban land use change with two explicit components: demand quantification and spatial allocation. The variation among models

largely stems from differences in conceptual and technical approaches to these two components. While many of the existing models appear to meet several of our criteria, none to our knowledge allow for a sufficient degree of non-stationarity in urban growth across both space and time. Moreover, many of the models have onerous input data requirements and computationally expensive and complex model calibration procedures. Therefore, we developed a new model (SPRAWL) that meets all of our design criteria; in particular, it consists of several interdependent components, including demand, allocation, and suitability modules loosely based on the CLUE-S ([Verburg et al., 2002](#)) and FORE-SCE ([Sohl, Saylor, Drummon, & Loveland, 2007](#)) models, but with the addition of a unique matching component that induces non-stationarity in urban growth across space and time, and the use of a novel state-space approach for spatially allocating growth to individual disturbance patches based on resistant kernels ([Compton, McGarigal, Cushman, & Gamble, 2007](#)) applied to suitability surfaces. We believe that these innovations provide a useful framework for modeling urban growth in other geographies and will serve to motivate the development of new and better ways to address non-stationarity in urban growth modeling systems.

There is recognition that the goal of urban growth modeling should not be accurate predictions of urban growth at the pixel level, but rather realistic patterns of growth overall at coarser scales as well as useful comparisons among alternative growth scenarios ([Li & Gong,](#)

2016). Moreover, the utility of urban growth modeling stems predominantly from understanding the consequences of land use change to inform land use planning and policy with regards to a wide variety of ecosystem services, such as the regulation of freshwater quantity and quality, sequestration of carbon, and the conservation of biodiversity (Verburg et al., 2004). In particular, planning for the conservation of biodiversity can be informed by scenario-based urban growth modeling, and there are numerous ongoing efforts to do so, for example those associated with the North American Landscape Conservation Cooperatives (Millard et al., 2012). As with the variety of urban growth models, there are myriad approaches for assessing the ecological impacts of urban growth. The DSL project includes a complementary two-pronged ecosystem- and species-based approach, each of which has multiple elements. Of particular interest here is the ecosystem-based assessment (McGarigal et al., 2018), which is based on the concept of “ecological integrity” defined as the ability of an area to support biodiversity and the ecosystem processes necessary to sustain biodiversity over the long term, especially in response to disturbance and stress. The ecological integrity assessment provides a basis for comparing the potential ecological impacts of alternative land use scenarios, which can then be used to inform landscape conservation design.

Our primary objective in this paper is to describe the SPRAWL model and its application to simulate land-use conversions associated with urban growth across the Northeast for 10-year time steps between 2010 and 2080 under a baseline scenario. Our secondary objective is to conduct a scenario-based analysis of the effects of urban growth on ecological integrity under scenarios of varying total demand and sprawl intensity, but here the focus is on illustrating the scenario-based utility of the SPRAWL model rather than the details of the ecological assessment.

## 2. Methods

We implemented our SPRAWL model in R (R Core Team., 2016) using custom packages for both parallel and raster processing. Fig. 1 outlines the structure of the SPRAWL model as described below.

### 2.1. Model training

#### 2.1.1. Training data

We selected three regions across the Northeast to be used as training areas: Maine, Massachusetts, and the Chesapeake Bay (Fig. 2). Maine and Massachusetts data came from the National Oceanic and Atmospheric Administration (NOAA) Coastal Services Center's Coastal Change Analysis Program (C-CAP) for the years 1996 and 2006. Chesapeake Bay data came from the Chesapeake Bay Watershed Land cover Data Series (CBLCD; Irani & Claggett, 2010) for the years 1984 and 2006, which was based partly on the C-CAP data as well. One strength of the matching approach described below is that the historical training data need not be available throughout the entire application region. However, it is important that the training data be representative of conditions throughout the application region and be directly comparable to the land-use layer for the initial (or starting) condition of the simulation. The three training regions were chosen because they each had a relatively long period of historical land-use GIS layers available that included multiple urban land-use classes (low-, medium-, and high-intensity development, and developed open space) mapped at the 30-m resolution. Together, they represented varied amounts of current land-use intensity and historical land-use conversion that likely represented the full gradient found across the Northeast. In addition, these data were directly comparable to the 2011 National Land Cover Data (NLCD; Homer et al. 2015) that is the most consistent and spatially comprehensive product of its kind currently available across the Northeast. We examined each of the developed classes in NLCD, C-CAP and CBLCD in comparison with aerial imagery and determined that the developed open space class was unreliable and inconsistently characterized, and it

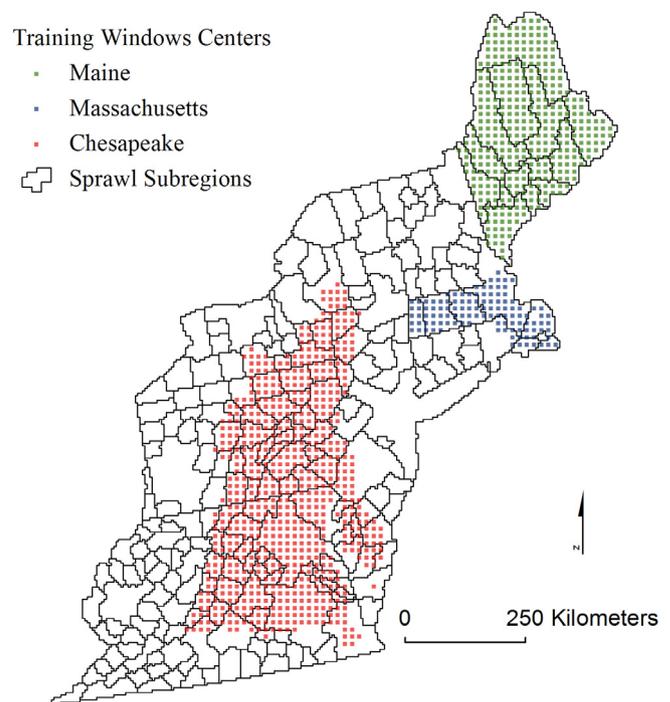


Fig. 2. Northeastern U.S. subdivided into 241 “subregions” defined by U.S. Census Bureau 2010 Core Base Statistical Areas (CBSAs) where they existed, or individual counties for those not in CBSAs, that we used for projecting the demand in urban growth forecasts, and the three areas (Maine, Massachusetts, and the Chesapeake Bay) used to train the SPRAWL model.

had the lowest producer’s accuracy of any of the developed categories in the NLCD (e.g., Wickham et al., 2017), so we opted to treat it as static in the urban growth model (i.e., the footprint of land classified as developed open space remains constant during an urban growth simulation). The other three developed land classes generally performed better in the NLCD (e.g., Wickham et al., 2017) and were fairly consistent between the data sources.

#### 2.1.2. Development transitions

For each training region, we compared the raster from the earlier time step to the later one to create a change raster for each of six transition types: undeveloped to each of 1) low-, 2) medium-, and 3) high-intensity development, and the three possible increases in development intensity: 4) low-to-medium, 5) low-to-high, and 6) medium-to-high. We did not consider transitions to lower urban intensities or transitions back to a non-urban or undeveloped condition (i.e., “rewilding”). While these transitions did occur, they were unreliably assessed in the training data and, moreover, were extremely rare; thus, we deemed these transitions of minor importance and omitted them from the urban growth model.

#### 2.1.3. Training windows

We divided each training region into non-overlapping square training “windows” approximately 15 km on a side (498 30-m cells, or 14.94 km on a side). The choice of window size reflected a compromise between keeping the windows as small as possible to reflect meaningful local variation in landscape context that might affect urban growth and keeping the windows large enough to avoid the idiosyncrasies of very small landscapes (and to reduce computer processing time). We used the change raster to calculate and store for later use in the matching algorithm (see below) three statistics for each training window:

- 1) *Rate of new urban development* — number of cells that underwent new development (i.e., transition types 1–3) during the training

period, divided by the length of the training period, and multiplied by 10 to convert to amount in cells per decade.

- 2) *Distribution of transition types* — proportion of the total number of cells exhibiting any transition (transition types 1–6) during the training period comprising each of the six transition types.
- 3) *Patch size distribution* — distribution of development patch sizes, where a patch was defined as a set of contiguous cells (based on an eight-neighbor rule) that underwent the same development transition. Note, because the Chesapeake Bay training data period was 22 years (instead of 10), there was undoubtedly a bias high in the computed patch size distribution for these training windows, since there was roughly twice the time for individual developments to coalesce into larger patches. While this bias was real, we deemed the overall effect on the coarse-scale patterns of urban growth to be minimal.

#### 2.1.4. Building suitability models

For each training window, we computed the Gaussian kernel (12.8 km bandwidth) intensities of adjusted development and of open water based on the land cover map developed for the DSL project (McGarigal et al., 2017), which was consistent with the training data sets with regards to the mapped developed land uses, and converted these values to z-scores (i.e., mean = 0 and standard deviation = 1). This allowed us to place each training window in a standardized two-dimensional state-space defined by the intensity of adjusted development and open water (Fig. 3), which we determined to be the two most important variables affecting allocation based on preliminary analyses. We defined adjusted development by assigning weights to cells as follows: 1 = low-intensity development, 2 = medium-intensity development, 3 = high-intensity development, n/a = “non-buildable” cells, and 0 = buildable cells. We treated cells classified as any land cover class associated with open water, wetland, roads and developed, in addition to cells designated as secured lands (i.e., already protected from future development; derived from The Nature Conservancy (2011) secured lands layer), as “non-buildable”, and all other cells as “buildable”.

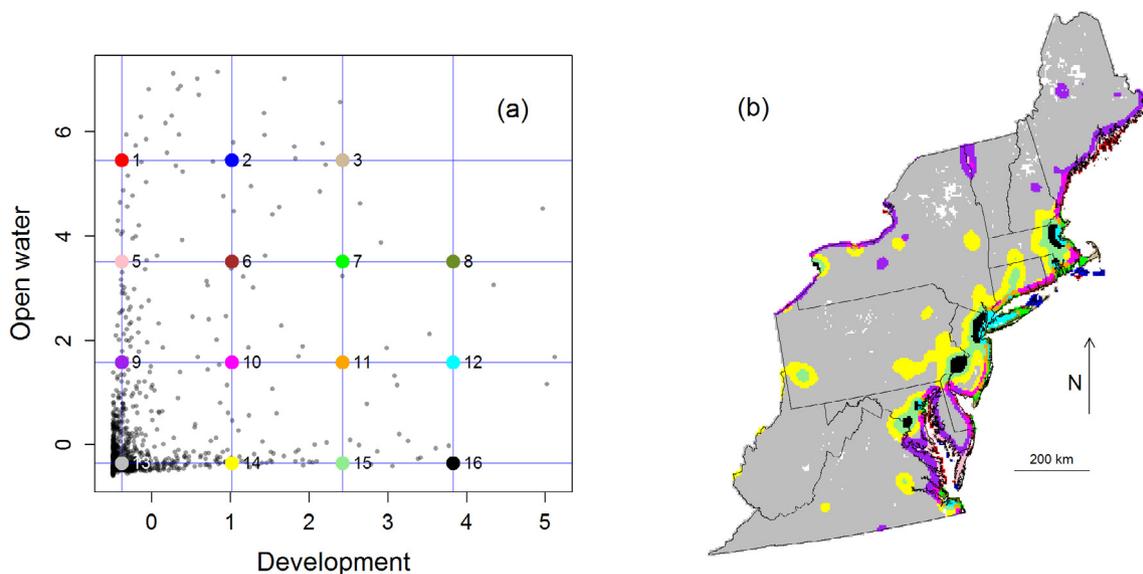
We located 16 uniformly distributed “model points” or locations in this state-space (Fig. 3). For each of these points, we fit separate binary logistic regression (i.e., “suitability”) models for each of the six transition types to a set of training points (i.e., randomly selected cells of

the corresponding transition type at least 150 m apart to avoid pseudo-replication, matched with an equal number of randomly selected buildable cells excluding cells that transitioned during the training period) from the training windows located within 1, 1.5 or 2 standard deviations (on the z-scored axes of the state space) from the model point in the state-space. We used the smallest of the three standard deviations needed to achieve a minimum of 200 (mean = 3,458; range = 208–60,449) training points (half of which experienced that transition and half of which were buildable points) in order to have the fitted model most closely reflect landscape conditions described by that model point in the state-space (Fig. A1). If we could not meet the minimum sample size of training points within 2 standard deviations of the model point, we dropped the model point from consideration.

For each model point and transition type, we fit all-subsets logistic regression models based on a set of 10 predictors, including the intensity (at multiple kernel bandwidths) of open water, primary and secondary roads, all roads except motorways, and adjusted development, in addition to transformed percent slope and distance to nearest road (Appendix A). This suite of predictor variables was informed by preliminary analyses that involved a comprehensive hierarchical model selection process including 103 different predictors representing combinations of a large number of candidate variables and wide range of spatial scales. Finally, we used full model averaging based on AIC model weights for models within 4 AIC units of the best model to predict each of the six transition types for each of 12–16 model points (depending on transition type) (see Appendix A for detailed results).

#### 2.2. Model execution: Simulating urban growth

The SPRAWL model involves simulating urban growth according to a user-specified scenario based on the algorithm described below. We use as the basis for the “current” or initial land-use condition the set of developed land cover classes in the DSL land cover raster which are based on 2011 NLCD, although any land cover classification would suffice so long as it is consistent with the training data and the derived suitability models. We subdivide the entire application region into non-overlapping square application “panes” ~5 km (166 cells, or 4980 m) on a side, each of which is embedded as the central pane within a square application “window” consisting of 3x3 panes (~15 km on a side, identical to the training windows). We implemented overlapping



**Fig. 3.** (a) Standardized two-dimensional model state-space defined by the intensity of adjusted development (see text for details) and open water within which each training window (~15 km on a side, depicted by gray points) was located and the 16 uniformly distributed model points for which we developed suitability models (see text for details); and (b) the corresponding spatial distribution (at the ~5 km application pane resolution) of the 16 model points depicted in the model state-space.

windows to prevent major boundary effects across application panes. In this way, each application pane shares 2/3 of its window with each of its four nearest neighboring panes, preventing abrupt differences from occurring between neighboring panes.

### 2.2.1. Determining demand

Prior to model execution, we establish the “demand” for additional cells of each urban land cover class in each 10-year time step. The demand dictates the overall amount (in cells) of urban land-uses to allocate throughout the area of interest in each time step. Although demand could be derived from any suitable data source, we chose to base demand on county-level forecasts derived for a U.S. Forest Service 2010 Resources Planning Act (RPA) assessment (Wear, 2011), after resolving some issues with the data (Appendix B). We assigned each application pane to a single county (or merged county) based on a majority rule, producing a gridded version of the RPA county map at 5 km resolution, and then aggregated the application panes into U.S. Census Bureau 2010 Core Base Statistical Areas (CBSAs) where they existed, or retained the forecasts at the county level for those counties not in CBSAs, resulting in a total of 241 geographic units, which we refer to as “subregions” (Fig. 2). We converted the RPA forecasts (which were given in absolute area) to development rates, computed as the projected development divided by the land area (as reported in the RPA assessment), and then converted this to an absolute demand (in cells) for each subregion in our gridded version by multiplying the forecasted rate of development by the count of land cells within the gridded subregion (Appendix B for link to final spatial data). Note, calculating development rates based on land area as reported in the RPA assessment corrected for the discrepancy in land area between RPA counties and our gridded version.

### 2.2.2. Matching to training windows

Based on the pre-determined total demand (above), we simulate urban growth for each time step as follows:

First, the application window for each pane is matched to the three most similar training windows (from any of the three training regions) based on geographic proximity and four landscape metrics, including the intensity (Gaussian kernel bandwidth = 12,800 m) of 1) adjusted development (as defined above), 2) all roads, except motorways, and 3) open water; and 4) density of roads (except motorways) within the window. The selection of the landscape metrics was informed by preliminary analyses that involved a comprehensive model selection process involving 50 different variables and scales to determine which combination of variables and scales had the greatest ability to predict the amount and pattern of future growth in each training window.

As noted previously, the reason matching occurs at the window scale, but then the allocation described below is applied at the pane level (which represents 1/9 of the window size), is to reduce boundary effects across application panes. For each application and training window, we compute the four landscape metrics listed above and convert the values to z-scores (mean = 0 and standard deviation = 1). We then compute a similarity score between each application window and every training window based on Euclidean distance in z-scores, such that it ranges from 1 if all four metrics are the same and 0 if all four metrics differ by 6 standard deviations. We then multiply these similarity scores by a geographic adjustment that applies a penalty to training windows based on how much farther they are from the application window than the closest training window. The adjustment is based on a Gaussian function such that it ranges from 1 at small deltas (i.e., the closest training window gets an adjustment of 1) to 0.5 at very large deltas (Appendix C). The resulting adjusted similarity scores are rank-ordered and the three most similar (weighted by proximity) training windows are matched for each application window.

The proximity adjustment reflects the lack of regionally available and consistent spatial data on local socio-economic drivers (e.g., zoning laws, median income, etc.) that influence land-use change that are not

incorporated into the model. We recognize that application windows are likely to grow more like nearby windows simply because they have similar local factors acting upon them. To partially account for this local variation, we added the local weighting component described above to the matching algorithm. In this way, an application window in Maine is more likely to be matched with training windows in Maine than training windows in Massachusetts or Chesapeake Bay, with the expectation that they will grow more like Maine has grown historically. Moreover, the proximity adjustment helps alleviate potential issues from matching to training windows from different datasets.

### 2.2.3. Allocating urban growth

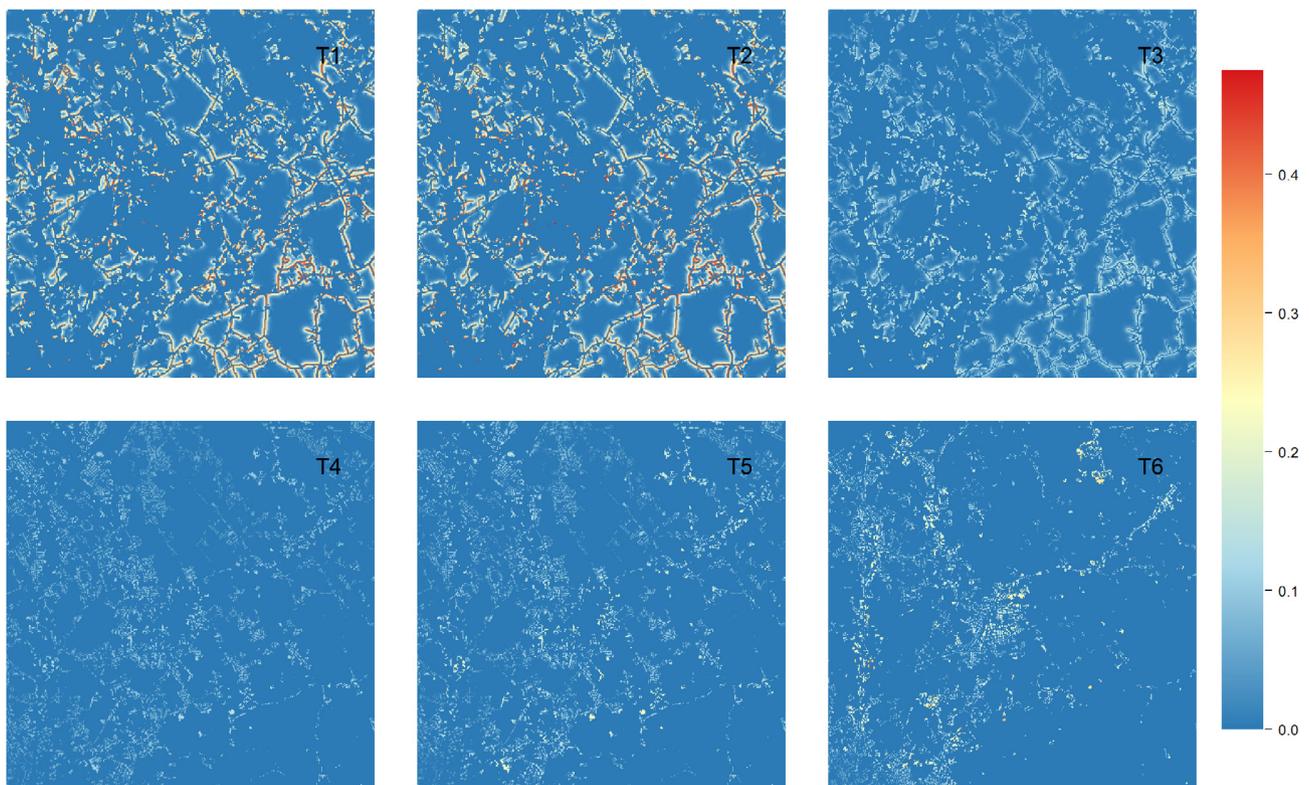
Once each application window is matched to three training windows, we allocate the total demand (in cells) for the current time step within each gridded subregion to each application pane. To do this, we calculate the average total amount of new development (i.e., transitions types 1–3) per decade observed historically in the three matched training windows. This “match amount” is multiplied by the proportion of the application window that is buildable and the proportion of buildable cells in the central pane. We also adjust this match amount as necessary to ensure that no more than 14% of the buildable cells for transitions 1–3 (i.e., available for new development) in the pane are built in any one decade. This maximum rate of development per decade was based on the 99<sup>th</sup> percentile of the corresponding distribution observed in the hindcast dataset we developed to validate the urban growth model (see below). The result is an interim measure of the amount to allocate to each application pane that reflects the historical distribution among panes having a similar landscape context. Lastly, the absolute demand (in cells) for each application pane in the current time step is computed by dividing the pane's interim match amount by the total interim match amount across all panes in the corresponding subregion and multiplying by the total demand in the subregion. In this manner, the total demand (in cells) for each subregion is allocated among application panes such that the more historical development that occurred in the matched training windows, the higher proportion of the future demand is assigned to the application pane.

Next, the demand (in cells) in each application pane for the current time step is allocated among the six transition types based on the historical distribution in the matched training windows, with the sum of the first three transitions (i.e., undeveloped to low-, medium- or high intensity developed) made to match the total allocation to the pane, and the ratio among all six transitions made to match the historical ratios in the matched training windows.

### 2.2.4. Determining suitability

Once the number of cells needed for each transition type in each application pane for the current time step is determined, we determine the probability of each transition type occurring at the cell level in each pane. To do this, for each transition type we create an inverse distance-weighted average logistic regression model based on the distance between the application window and each model point in the model state-space developed during the Training phase (Fig. 3). For this calculation, we force the distance to be at least 0.05 standard deviations for all model points so that a single logistic regression model cannot get 100% of the model weight. Next, we use these weighted-average models to compute the relative probability (i.e., suitability) of each transition type for each cell in the application pane (Fig. 4). Lastly, for transition types 1–3, we assign a zero value to all non-buildable cells. For transition types 4–6, we assign a zero value to all cells not of the focal class (e.g., for transition type 4, only cells of low-intensity development are allowed to have a non-zero value).

Importantly, because the Gaussian kernel (12.8 km bandwidth) intensity of adjusted development surface is changing over time due to urban growth, the position of each application window in the two-dimensional model state-space is shifting over time as well. Consequently, the patterns of urban growth in an application window will shift over



**Fig. 4.** Probability of development (i.e., suitability) at the 30 m cell level for a random application window and time step within the northeastern U.S. for each of six transition types (see text for details).

time and become more like the patterns characteristic of increasingly urbanized windows (see Results). This feature ensures non-stationarity in the patterns of urban growth over time, which we deem an important and distinctive feature of our urban growth model.

#### 2.2.5. Building development patches

Given the demand (in cells) for each transition type allocated to each application pane for the current time step and the corresponding suitability surfaces, we build development patches for each transition type, as follows:

- 1) First, randomly select a cell to initiate a disturbance based on the relative probability (i.e., suitability) surface for that transition, but if the candidate pool of buildable cells vastly exceeds the development target ( $> 500$  times as many cells in the top 15% than needed for the transition), we also limit the selection of cells to those in the top 15<sup>th</sup> percentile of suitability within the application pane; thus, cells with the highest suitability values are more likely to initiate a new development. Note, regardless of whether the selection is limited to the top 15<sup>th</sup> percentile of suitability values, cells are still selected based on their relative probabilities;
- 2) Next, randomly draw a patch size from the observed distribution of patch sizes in the three matched training windows for the corresponding transition type;
- 3) Next, spread outward from the initiation cell with a resistant kernel (Appendix D) to the desired patch size where resistance is based on the complement of the probability of transition. This produces patches that extend farther into neighboring cells that are more suitable for the transition.
- 4) Lastly, repeat the process above, building development patches sequentially, until the total allocation of cells for the transition type is exhausted in the application pane (Fig. 5).

#### 2.3. Model evaluation

It was impractical to truly validate the SPRAWL model as applied here because the outcome depended strongly on the RPA forecasts of future development which were based on several assumptions about global and U.S. socioeconomic and climate trends likely to affect future U.S. resource conditions and trends as reflected in the IPCC 4th Assessment A1B, A2, and B2 scenarios (IPCC, 2007), none of which can be verified without waiting to see if the predictions come true. However, hindcasting provided an alternative for evaluating model performance, which we describe in detail along with the results in Appendix E. Briefly, we used the NLCD change raster for the period 2001–2011, in which each cell was coded as either unchanged or one of our six transition types, to create an approximate year 2000 version of our DSL land cover raster by applying the same process to the 2001 NLCD raster that we used to create DSL land cover (McGarigal et al., 2017). Next, we used the DSL 2000 land cover raster as the initial condition and ran the SPRAWL model for a single 10-year time step to generate the probability of development (or suitability) surface for each transition type, but we set the demand for new development (transition types 1–3) within each subregion to be equal exactly to what we observed for the buildable portion of the landscape in the NLCD change raster, effectively removing demand from the model evaluation (since our purpose was not to evaluate the RPA forecasts).

To evaluate the performance of the SPRAWL model for the 2000–2010 time step, we computed four different statistics. First, we compared the average derived probability of development at cells that actually underwent a development transition in the NLCD change raster to that of available cells (and reported the ratio) for each of the transition types. Second, despite the fact that the data were appropriately considered observations of “presence-vs-available” rather than “presence-vs-absence”, since the cells that underwent a transition were simply a subset of those that could have undergone a transition given the same suitability, we computed the area under the Receiver

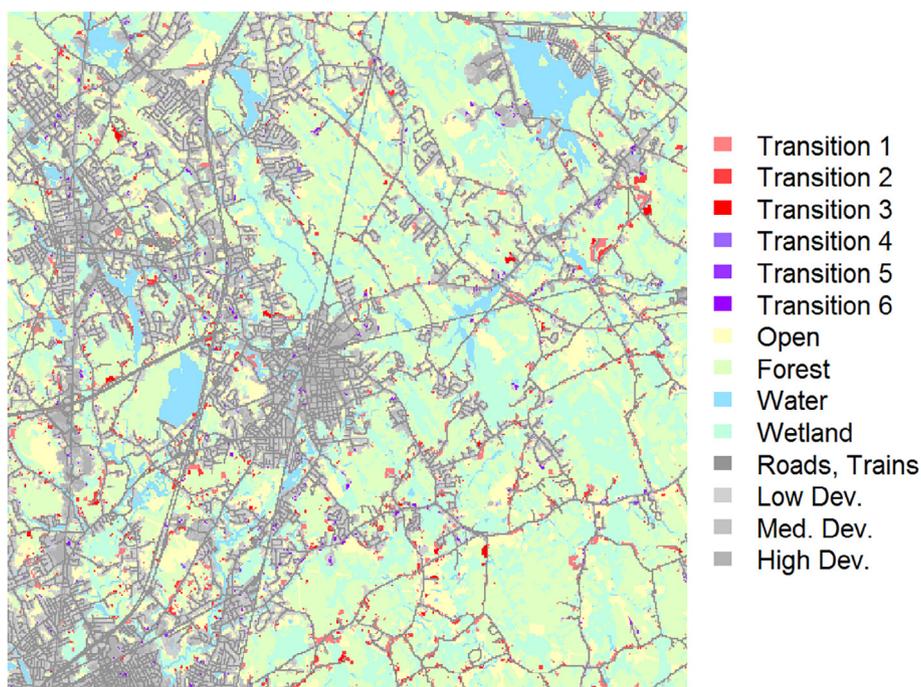


Fig. 5. Simulated development transition patches for the random application window and time step depicted in Fig. 4.

Operating Characteristic Curve (or AUC, e.g., Fielding & Bell, 1997) to gauge the model's discriminatory ability. In our context, AUC is equal to the probability that our model-derived probability of development ranked a randomly chosen transitioned cell higher than a randomly chosen available cell. Third, we computed the coefficient of concordance (Lin, 1989) between observed and predicted based on the methods recommended by Johnson, Nielsen, Merrill, McDonald, and Boyce (2006) for “presence-vs-available” data. The concordance coefficient measures the strength of agreement (in both accuracy and precision) between these two allocations based on their deviation from a 45-degree line (i.e., a line originating at 0 with a slope of 1), with values closer to 1 indicating better agreement and thus better overall model performance. Lastly, we computed the weighted skewness statistic based on the method recommended by Gregr and Trites (2008) for “presence-vs-available” data. This statistic measures the (weighted) skewness of the distribution of model-derived predicted probability of transition values, whereby the more predictive the model is the higher the proportion of observed transition points (after weighting to reflect availability) that are located at the higher end of the probability scale, resulting in a more left-skewed distribution.

#### 2.4. Model application: Ecological impacts of alternative urban growth scenarios

The SPRAWL model has two basic mechanisms for creating urban growth scenarios, both of which involve subjective deviations from the empirically-based baseline scenario. First, the overall amount of land that is developed (or the corresponding rate of development) can be modified from the baseline demand, which for this application we derived from the RPA assessment. Here we simply multiply the baseline demand by a factor to either increase or decrease the amount of development relative to the baseline; e.g., to increase the demand by 25% we would multiply the baseline demand by 1.25. In this scenario, the relative allocation of development among application panes and spatial patterns of development at the cell level would remain the same as the baseline, but the total amount of development would be 25% greater.

The other mechanism for creating scenarios is the “sprawl dial”, which allows us to create scenarios that are more “sprawly” than the

baseline. The baseline scenario acts as the status quo (sprawl level = 0) and distributes the number of new urban cells to each application pane in accordance with historical patterns using the allocation process described above. Thus, the baseline scenario exhibits sprawl, but at the level observed in the historical training data. Increasing the sprawl dial redistributes more of the allocation to panes further away from urban centers. Briefly, we create a variable-bandwidth Gaussian smooth of the baseline allocation, where the bandwidth ranges from 10 to 45 km, with larger bandwidths chosen when the focal pane is in proximity to a larger urban center (based on the 12.8 km Gaussian kernel intensity of adjusted development, as defined earlier). Thus, the development allocated to each baseline application pane is reallocated to itself and nearby panes based on the Gaussian kernel smooth, whereby the reallocation is distributed more broadly the larger the urban center. Generally the redistribution is forced further from the urban center, because the limited buildable area closer to the urban center restricts how much of the reallocation can be accommodated — resulting in the desired sprawl effect. For a relatively small urban center, this translates into redistributing the allocation over roughly a 10–20 km radius, and for a larger urban metropolitan area (e.g., Boston, Philadelphia, New York, Washington DC, etc.) the reallocation occurs over roughly 45–90 km. Thus, for the same setting of our sprawl dial, the realized scale of sprawl varies with the size of the urban center. Finally, we compute the weighted average of the standard baseline allocation and the smoothed allocation, with the weight determined by the setting of the sprawl dial. Thus, as the sprawl dial is increased between 0 and 1, more of the allocation to urban areas is dispersed to the outlying suburban areas based on the smoothed allocation, which determines how contagious (compact vs. “sprawly”) growth patterns will be at the broad scale compared to the baseline.

For this application, we analyzed four scenarios in a two-way factorial combination of two levels of demand (baseline and 25% increase) and two levels of sprawl (baseline and sprawl level = 0.75). See Appendix F for images depicting each of these scenarios. For each scenario, we ran three replicate 70-year simulations (2010–2080). For each simulation, we computed the *Index of Ecological Impact* (*ecoImpact*, McGarigal et al., 2018) for the year 2080 (Appendix G). Briefly, *ecoImpact* is derived from a composite of 6–18 ecological integrity

**Table 1**

SPRAWL model evaluation based on hindcasting, including for each of six development transition types (1–3 represent new development from undeveloped land; 4–6 represent redevelopment) the following statistics: 1) ratio of the average computed probability of development for cells that underwent development in the hindcast dataset (~2000–2010) to that of the corresponding available cells, 2) area under the Receiver Operating Characteristic Curve Area (AUC), 3) coefficient of concordance (CC), and 4) weighted skewness statistic.

Transition type	Use/availability ratio	AUC	CC	Weighted skewness
1 (undeveloped to low-intensity developed)	11.14	0.93	0.70	−0.73
2 (undeveloped to medium-intensity developed)	11.70	0.94	0.77	−0.71
3 (undeveloped to high-intensity developed)	9.86	0.93	0.80	−0.61
4 (low- to medium-intensity developed)	1.08	0.53	0.93	0.05
5 (low- to high-intensity developed)	1.25	0.60	0.98	−0.31
6 (medium- to high-intensity developed)	1.26	0.57	0.77	−0.17

metrics, depending on the ecosystem, measuring different aspects of ecosystem intactness and resiliency and provides an index of the change in ecological integrity between the current and future time steps relative to the initial condition. A site that experiences a major loss of ecological integrity due to urban growth has a high predicted ecological impact. *ecoImpact* is a raster, but here we sum the raster and average across the three replicate simulations for each scenario to produce a scalar index for comparison among scenarios.

### 3. Results

**Model evaluation.**—The average probability of transition derived from the SPRAWL model applied to the hindcast dataset was roughly 10–11 times greater for cells that realized new development (transition types 1–3) between 2000 and 2010 than that of available cells (i.e., all undeveloped buildable cells), indicating that the SPRAWL model was very effective in discriminating undeveloped sites likely to transition to new low-, medium-, or high-intensity development. This was reflected in the relatively high AUC values for new development and the relatively large negative values of the weighted skewness statistic (Table 1), which can also be visualized graphically (Appendix E). However, the SPRAWL model was weakly discriminating and therefore had weak predictive ability for redevelopment transitions (i.e., transitions from already developed to a higher-intensity development,  $AUC < 0.6$ ), which was not surprising given the relatively poor ability of the NLCD map to accurately detect redevelopment changes (Wickham et al., 2017). The SPRAWL model was also moderately well calibrated for the Northeast landscape, with coefficients of concordance between observed and predicted  $> 0.7$  for all 6 transition types (Table 1), which can also be visualized graphically (Appendix E).

**Non-stationarity.**—As intended, the SPRAWL model exhibited spatial and temporal non-stationarity in the allocation of future development (based on the specified demand to each subregion) to individual 5 km application panes and in the distribution to individual development patches within panes based on the derived suitability surfaces (Fig. 6). The shifting spatial allocation of new development over time reflected the evolving development of a subregion, such that as panes were developed they “behaved” more like similarly developed panes in the historical training data. Differences between the two development

scenarios illustrated in Fig. 6 reveals that in the increased demand and sprawl scenario more of the reallocation of future development was shifted away from the urban center to the outskirts of the subregion, particularly to the northwest in what might be considered a “sprawl zone”.

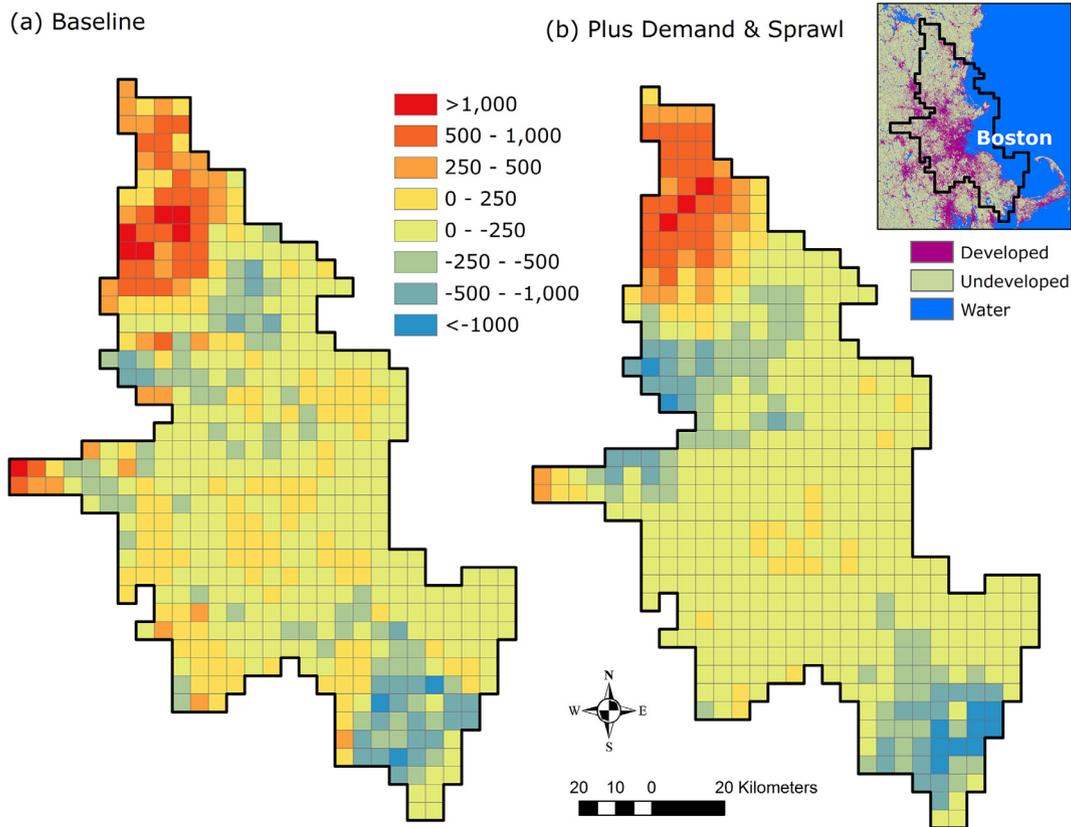
As intended, the model also exhibited non-stationarity in the cell or patch-level distribution of transitions. Application panes occupied a position in the two-dimensional model state-space in 2010 based on intensity of open water and development in the application window they were embedded within. The position of the pane in this state-space determined the weighted-average logistic regression model used to create the suitability (i.e., probability of development) surface, which stochastically determined where new development patches were placed. As the pane was developed over time, its position in the state-space shifted along the x-axis (Fig. 7), resulting in an evolving suitability surface. Differences between the two development scenarios illustrated in Fig. 7 was reflected in the longer arrows (i.e., greater shift in the state space over time) in the increased demand and sprawl scenario.

**Model application: scenario-based ecological assessment.**—The total *ecoImpact* under the baseline urban growth scenario varied little among replicate simulations (coefficient of variation = 0.2%), and the replications varied even less ( $CV < 0.02\%$ ) within each of the three alternative scenarios, suggesting that a single simulation for each scenario was sufficient to describe the overall ecological impact when pooled across the entire regional landscape. The 25% increase in demand scenario resulted in an ~11% increase in total *ecoImpact* compared to the baseline scenario, indicating that there was not a 1:1 relationship between development and impact. Indeed, while some locations experienced relatively high impact, some of the additional development occurred in places that suffered little ecological impact (i.e., infilling of already developed areas). The increased sprawl scenario resulted in a ~3% increase in total *ecoImpact* compared to the baseline scenario, indicating that although greater sprawl than was observed under the baseline scenario (reflecting historical trends) had a negative ecological impact, the magnitude of impact was relatively small for the increased level of sprawl that we simulated. Additional scenarios are needed to fully examine how *ecoImpact* varies as a function of greater levels of sprawl, and further research is needed to calibrate our sprawl dial with other established metrics for quantifying sprawl (e.g., Frenkel & Ashkenazi, 2008; Reis, Silva, & Pinho, 2015; Sudhira, Ramachandra, & Jagadish, 2004). Lastly, the increased demand and sprawl scenario resulted in a ~14% increase in total *ecoImpact* compared to the baseline scenario, indicating that the impact of increased demand and increased sprawl was largely additive.

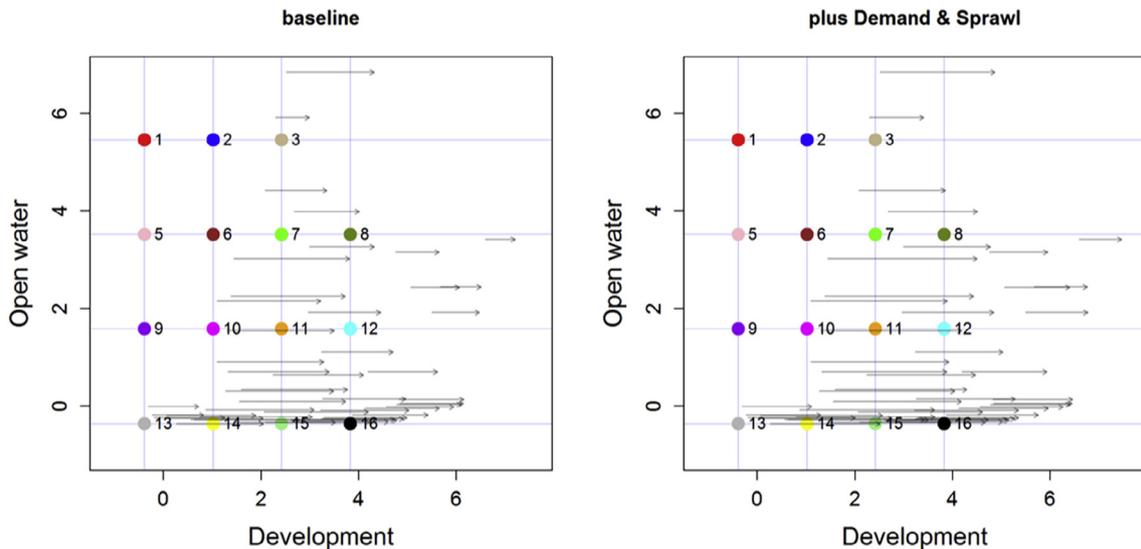
### 4. Discussion

The staggering variety of extant urban growth models (Li & Gong, 2016) undoubtedly stems from the diversity of application-specific modeling objectives coupled with practical constraints imposed by things such as data availability and computational resources. We developed the SPRAWL model to meet several design criteria specific to its application within the broader modeling framework of the Designing Sustainable Landscapes project and with respect to the following important considerations.

**Addressing demand for urban growth.**—One of the biggest challenges in urban growth modeling is determining the future demand for urban land use. Not surprisingly, one of the major distinctions among urban growth models is related to methods for addressing future demand (Li & Gong, 2016). For example, agent-based models (ABMs) approach demand as an emergent property of the modeled behaviors or decisions of individual agents (Crooks, Castle, & Batty, 2008; Matthews et al., 2007; Parker et al., 2003); thus, demand is an endogenous process in the model. Models variously classified as cellular automata (CA, Santé et al., 2010), on the other hand, often generate demand via projecting



**Fig. 6.** Heat maps depicting the difference between the 2080 and 2010 allocation to the application panes (~5 km on a side) within the Boston metropolitan subregion in the northeastern U.S. under two urban growth scenarios: (a) baseline and (b) increased demand (25% greater than baseline) and sprawl (sprawl level = 0.750). Warmer colors are associated with positive deltas (i.e., panes receiving more allocation of new development in 2080 compared to 2020); conversely, cooler colors are associated with negative deltas (i.e., panes receiving less allocation of new development in 2080 compared to 2020). Deltas are computed as the difference in the number of cells of new development (i.e., transitions 1–3, see text for details) between 2020 and 2080. Non-zero deltas indicate non-stationarity; i.e., new development was allocated among panes differently in 2080 compared to 2020.



**Fig. 7.** Illustration of non-stationarity in the suitability surfaces for 50 randomly selected application panes (~5 km on a side) within the Boston metropolitan subregion in the northeastern U.S. under two urban growth scenarios: (1) baseline and (2) increased demand (25% greater than baseline) and sprawl (sprawl level = 0.750). Plot depicts the trajectory of change in the position of the selected application panes in the model state-space between 2010 and 2080 for transition 1 (new low-intensity development). Each pane moves through state-space on the x-axis (intensity of development) to a varying degree depending on how much development it receives. The varying magnitude of change among panes, as reflected by the varying length of the vectors, represents non-stationarity in urban growth. The vectors are longer in the increased demand and sprawl scenario because of the greater amount of development simulated. Any movement through state-space over time results in a change in the derived suitability surface that in turn affects the stochastic development, indicating non-stationarity in urban growth.

forward statistical trends derived from the historical training data (e.g., [Sohl and Sayer, 2008](#)). For the SPRAWL model, we chose instead to rely on the expertise of others and treat demand as an entirely exogenous process, whereby we simply input the demand as a constraint on urban growth. For this application, we used the RPA forecasts developed by [Wear \(2011\)](#) which combined historical trends with econometric projections of future human population and land use tied to the SRES scenarios ([Nakicenovic, Alcamo, & Davis, 2000](#)), which we deemed to be more robust than simply extrapolating the historical trends in our training data. In addition, the RPA forecasts exist at the county level, which we deemed a sufficiently fine spatial resolution to discern important spatial patterns across the Northeast, and the RPA assessment is done every 10 years which provides a consistent basis for updating demand in future urban growth applications.

*Spatial allocation of urban growth.*—Another challenge that lies at the core of urban growth modeling is determining how to spatially allocate demand. There are perhaps as many different approaches as there are models. Most approaches involve transitioning cells or patches typically within some higher level spatial unit (e.g., county, land cover type, ecoregion) based on some form of a suitability surface or set of transition rules applied at the cell level until the demand is met (e.g., [Chaudhuri & Clarke, 2013](#); [Sohl et al., 2007](#); [Theobald 2005](#); [Veldkamp & Fresco 1996](#); [Verburg, DeKoning, Kok, Veldkamp, & Bouma 1999](#); [White & Engelen 2000](#); [Zhao & Murayama 2007](#)). In the SPRAWL model, we implement a unique multi-level allocation of growth within subregions (CBSAs or counties) that involves first allocating growth to overlapping application panes (5 km on a side) and then to individual disturbance patches within panes. The pane-level allocation is based on a unique matching algorithm that matches each application window to similar windows in the historical training data based on several landscape metrics describing the intensity of development, roads and open water in addition to geographic proximity. Thus, each pane receives its proportional allocation of each urban transition type based on how historical growth was allocated to similar landscape contexts. The allocation to each pane is then distributed among individual disturbance patches based on the derived suitability surface for that landscape context. In this manner, the cell or patch-level disturbance patterns emulate the fine-scale patterns of growth observed historically in similar landscape contexts, at least to the extent that they are controlled by factors such as the intensity of open water and development at multiple scales, slope and proximity to roads. Moreover, similar to some other models (e.g., [Chen, Li, Liu, & Ai, 2014](#); [Meentemeyer et al., 2013](#); [Sohl & Sayer, 2008](#)) we transition patches of cells rather than individual cells independently and do so organically by selecting initiation cells and growing outward based on the suitability surface using a unique resistant kernel algorithm ([Appendix D](#)) until a randomly selected patch size is met based on the historical distribution of patch sizes in the matched training windows. Importantly, this multi-level and multi-scale allocation is computationally efficient and requires no model calibration. Our model evaluation based on hindcasting suggests that this approach was very effective in representing the recent historical urban growth patterns.

*Local suitability.*—Most urban growth models have some form of suitability module for deploying urban growth at the cell or patch level, and as with allocation, there again appear to be almost as many approaches as there are models. After evaluating several alternative approaches, we opted to use logistic regression based on a parsimonious suite of predictors at multiple scales that are easily computed from a land cover map, similar to other models (e.g., [Hu & Lo, 2007](#); [Sohl & Sayer, 2008](#); [Wu, 2002](#)). However, we recognized that the factors that influenced urban growth patterns at the cell or patch level were not necessarily the same in all landscape contexts; in particular, the factors that drove growth in urban areas and near the coast and other large water bodies differed from those in more rural and terrestrial settings. To address this variability we developed the suitability model state-space approach, in which we used the training data to develop model-

averaged logistic regression models for “points” distributed evenly throughout a two-dimensional state-space defined by the intensity of development and open water. This state-space captures two dominant gradients affecting local urban growth patterns in the Northeast, but it could be based on other user-selected gradients for other applications. Once the state-space is defined and logistic regression models are built for each model point and transition type, during the urban growth simulation each cell, as part of an application window, is positioned in the state-space and an inverse-distance weighted logistic regression model is applied. In this manner, the suitability surface derived from the logistic regression model varies as an application window moves through the state-space. This is analogous to geographically-weighted regression, but instead of geographic space the model varies in ecological (i.e., land cover) space.

*Treatment of scale.*—The importance of explicitly dealing with scaling issues in urban growth models is generally recognized ([Verburg et al., 2004](#)), and increasingly models are adopting a multi-scale approach in either the structure of the model or in the quantification of the driving variables or both ([Li & Gong, 2016](#)). The multi-scale approach acknowledges that different driving forces behind land use change are important at different scales, but there are myriad ways of dealing with it in a model. In the SPRAWL model, we chose to incorporate multiple scales into both the multi-level structure of the model and the quantification of the driving variables. Regarding the multi-level structure of the model, demand is input at the subregional level (CBSA or county), but then allocated to the application pane level, and then to the individual cell or patch level. This hierarchical three-level structure is implicitly multi-scale as well, as subregions encompass larger spatial extents than application panes (5 km on a side), which are larger than individual cells (30 m) or disturbance patches. In addition, the landscape variables that drive the pane-level allocation of demand in the matching algorithm and that drive the suitability surfaces within each pane can be quantified at multiple scales via kernel smoothing of the predictors at different bandwidths. Ultimately, the choice of scales for the subregions and panes and the bandwidths for the variables in the matching and suitability modules provides a flexible framework for incorporating the multi-scale nature of land use change.

*Mechanisms governing urban growth.*—Urban growth models typically are a hybrid of empirically-driven (statistical) and process-based (mechanistic) approaches ([Li & Gong, 2016](#)). However, due to the complexity of the mechanisms governing urban growth (at multiple scales) and the challenges in parameterizing such a model, we opted for an entirely data-driven approach in the SPRAWL model. Thus, the allocation of demand to application panes and then to individual cells and patches is driven by statistical relationships observed between urban development and various explanatory variables in the historical training data, without any attempt to infer causal relationships. While the use of correlative relationships in the model may be deemed a limitation, it allows the model to be computationally efficient and avoids the need for a comprehensive theoretical understanding of the mechanisms that drive urban growth and the challenging parameterization of the corresponding model.

*Data and computing requirements.*—An important consideration in any modeling application is the data and computing requirements, and urban growth models vary markedly in this regard ([Li & Gong, 2016](#)). We sought a model that could be implemented with minimal data requirements and efficiently at the spatial resolution of 30 m cells over the extent of the northeastern U.S. (64.5 million ha). Accordingly, we sought to avoid a computationally expensive internal model calibration process that is common in many models (e.g., SLEUTH, [Chaudhuri & Clarke, 2013](#)). In addition, NLCD is the only land cover product available across the entire Northeast with multiple urban intensity classes (i.e., low, medium, and high). However, we sought to use training data that were more accurate than NLCD, even if they were incomplete for the entire Northeast so long as they were representative of land use patterns observed throughout the region. We developed the unique

matching algorithm in the SPRAWL model in part for this purpose, which allows growth to be allocated locally to any application pane outside of the training regions based on similarity to matching windows in the training regions. This afforded us the opportunity to use the NOAA C-CAP and CBLCD data sources that were more accurate than NLCD (e.g., C-CAP, 2014) but there were available only for a portion of our entire application region (Fig. 2). Note, we retained the complete NLCD 2001–2011 change data as a somewhat independent, albeit less accurate, dataset for model evaluation via hindcasting.

*Non-stationarity in urban growth.*—One of the biggest challenges confronting urban growth modeling is dealing with non-stationarity in urban growth over both time and space (Li & Gong, 2016). In this context, non-stationarity refers to variation over time and/or space in the functional drivers of land use change, which is largely manifested in the treatment of model parameters that control urban growth. Indeed, this has proven to be quite challenging owing to our limited understanding of land use change drivers, but perhaps more so because of unpredictable changes in socioeconomic and political conditions that drive land use decisions and policies. Consequently, most models either assume stationarity (e.g., SLEUTH, Chaudhuri & Clarke, 2013) or incorporate a limited degree of non-stationarity (e.g., SERGoM, Theobald, 2005). Perhaps the most distinctive feature of the SPRAWL model is the degree to which we have induced non-stationarity in urban growth. First, as with many urban growth models, the user-specified exogenous demand allows for non-stationarity over time and among subregions via the temporally and spatially explicit forecasts. This feature allows subregions to vary over time in response to a non-linear forecast in demand (as is the case in the RPA forecasts we used) as well as allowing the allocation among subregions to vary over time (as is also the case in the RPA forecasts). Second, our unique matching algorithm induces non-stationarity over time and space at the application pane level, such that individual panes receive more or less of the subregional demand over time as their landscape character changes, as shown in Fig. 6. Lastly, our state-space suitability algorithm induces non-stationarity over time and space at the cell and patch level because the averaged logistic regression models for predicting relative suitability change as the application pane shifts its position in the state-space over time, as shown in Fig. 7.

*Scenario-based analysis.*—The utility of scenario-based analysis of urban growth simulations is widely accepted (e.g., Li & Gong, 2016) and there are numerous examples of using scenarios to evaluate land use change under alternative futures (e.g., Berling-Wolff & Wu, 2004b; Jantz, Goetz, & Shelley, 2003; Xiang & Clarke, 2003; to name just a few). Increasingly, scenario-based analysis is being used to evaluate the ecological impacts of alternative land use futures for the purpose of conservation planning (e.g., Eppink, van den Bergh, & Rietveld 2004; Seto, Guneralp, & Hutyra 2012; Sim & Mesev 2014). Here, for the main purpose of demonstrating the utility of the SPRAWL model for scenario-based ecological assessment and conservation planning, we conducted a limited scenario analysis to evaluate the effect of increasing the demand and the “sprawlness” of future development on ecological integrity in the Northeast. We demonstrated that greater sprawl for the same total amount of development has a negative effect on ecological integrity and, not surprisingly, that ecological impact increases as demand for new development increases. Additional investigation is needed to determine the shape of these relationships as both demand and sprawl increase even further. In addition, more comprehensive scenario-based analyses of the ecological impacts of urban growth under alternative landscape conservation designs is being done as part of the overall DSL project and will be reported on separately.

## 5. Conclusions

The SPRAWL model we presented here is one component of a more comprehensive landscape change, assessment and design (LCAD) modeling system that we developed for the DSL project. Importantly,

the model has a multi-level and multi-scale structure that aims to capture urban growth drivers operating at multiple scales. Due to this structure and the unique matching and state-space suitability algorithms, the model induces a relatively high degree of non-stationarity in urban growth across both time and space. In addition, the model is fully empirically-based, obviating the need for understanding and parameterizing a complex process-based model, and has minimal input data requirements.

Of course no model is without limitations (e.g., quality of training data), and the SPRAWL model is no exception. First, the model's dependence on the historical pattern means that any errors or biases in the historical land use mapping may be reproduced by the model. We detected one such bias (underrepresentation of development near roads) and partially corrected it (see Appendix A), but others may exist. There is a clear need for more accurate region-wide mapping of development intensity.

Second, the exogenous treatment of demand means that the burden of addressing socioeconomic drivers of urban growth is deferred to experts in the field. Indeed the SPRAWL model has no endogenous mechanisms for explicitly considering socioeconomic drivers, other than through their correlation with land cover characteristics as represented in the historical training data. In the Northeast, in particular, we recognize that local socioeconomic factors such as local zoning bylaws play an important role in determining where urban growth can and cannot occur. Unfortunately the spatial data for these factors across the entire Northeast do not yet exist. We tried to indirectly incorporate these effects, but only in part and at a relatively coarse scale, by adding the geographic distance between application window and training windows to the matching algorithm, but ultimately this is not a satisfactory solution.

Third, roads are undeniably a hugely important driver of new development. Currently, common to most urban growth models, the SPRAWL model has no provision for building new roads, or expanding existing roads to support or direct future urban growth. Roads are disproportionately important in generating the suitability surfaces, and generally constrain new development to occur very close to existing roads. Thus, the model is biased against the building of subdivisions into currently non-roaded areas. While this is an important limitation of the model, we could not devise a justifiable method for predicting new road development, so this remains an important subject for future research. Similarly the model has no provision for “rewilding” or converting from a developed class to an undeveloped class owing to limitations in the historical land cover data; however, this shortcoming is easily rectified should more reliable land cover data become available in the future.

Lastly, the SPRAWL model operates under the general assumption that the factors that affected the local amount and spatial pattern of development in the recent past (as reflected in our allocation of demand to individual application panes and the corresponding suitability surfaces) will persist into the future. Is the past a reliable guide for the future? This question is at the heart of a major challenge for most empirically-based urban growth models. Without a detailed and accurate mechanistic understanding of the socioeconomic and political forces that drive land use decisions at multiple scales, it would seem to us that using the past is perhaps our best option, as long as we continually update our inferences based on the most recent trends and predictions of altered resource utilization patterns due to factors such as climate change (Jones, Watson, Possingham, & Klein, 2016).

Despite the limitations above, we posit that the SPRAWL model can simulate realistic enough future land use patterns that when combined with meaningful ecological assessments can help provide useful guidance to landscape conservation design.

## Acknowledgements

This work was supported by the United States Fish and Wildlife

Service, North Atlantic Landscape Conservation Cooperative (NALCC), United States Geological Survey Northeast Climate Science Center, and the University of Massachusetts, Amherst, United States. We especially thank Andrew Milliken and Scott Schwenk of the NALCC for their continued support and close involvement in several conservation applications involving the DSL project and the use of the SPRAWL model.

## Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.landurbplan.2018.04.018>.

## References

- Allen, C. R. (2006). Sprawl and the resilience of humans and nature: an introduction to the special feature. *Ecology and Society*, 11(36) URL: <http://www.ecologyandsociety.org/vol11/iss1/art36/> [Date accessed: March 1, 2017].
- Berling-Wolff, S., & Wu, J. (2004). Modeling urban landscape dynamics: a review. *Ecological Research*, 19, 119–129.
- Berling-Wolff, S., & Wu, J. (2004). Modeling urban landscape dynamics: a case study in Phoenix. *Urban Ecosystems*, 7, 215–240.
- Chaudhuri, G., & Clarke, K. C. (2013). The SLEUTH land use change model: a review. *The International Journal of Environmental Resources Research*, 1, 88–104.
- Chen, Y., Li, X., Liu, X., & Ai, B. (2014). Modeling urban land-use dynamics in a fast developing city using the modified logistic cellular automaton with a patch-based simulation strategy. *International Journal of Geographic Information Science*, 28, 234–255.
- C-CAP. 2014. National 2010 coastal change analysis program accuracy assessment report. National Oceanic and Atmospheric Administration, Office for Coastal Management, September 2014. (online) URL: <https://coast.noaa.gov/data/digitalcoast/pdf/ccap-assessment-national.pdf> (Date accessed: 1.03.17).
- C-CAP Land Cover Atlas. (online) URL: <https://coast.noaa.gov/digitalcoast/tools/lca> (Date accessed: 25.03.11).
- Chesapeake Bay Land Cover Data Series (CBLCD). (online) URL: [http://landsat.gsfc.nasa.gov/pdf\\_archive/CBWLCD.pdf](http://landsat.gsfc.nasa.gov/pdf_archive/CBWLCD.pdf) (Date accessed: 25.03.11).
- Compton, B. W., McGarigal, K., Cushman, S. A., & Gamble, L. R. (2007). A resistant-kernel model of connectivity for amphibians that breed in vernal pools. *Conservation Biology*, 21, 788–799.
- Crooks, A., Castle, C., & Batty, M. (2008). Key challenges in agent-based modeling for geospatial simulation. *Computers, Environment and Urban Systems*, 32, 417–430.
- DeNormandie, J., & Corcoran, C. (2009). Losing ground beyond the footprint: patterns of development and their impact on the nature of Massachusetts. Fourth Edition of the Losing Ground Series, MassAudubon. [online] URL: [http://www.massaudubon.org/content/download/8601/149722/file/LosingGround\\_print.pdf](http://www.massaudubon.org/content/download/8601/149722/file/LosingGround_print.pdf) (Date accessed: 1.03.17).
- Eppink, F. V., van den Bergh, J. C. J. M., & Rietveld, P. (2004). Modeling biodiversity and land use: urban growth, agriculture and nature in a wetland area. *Ecological Economics*, 51, 201–216.
- Fielding, A. H., & Bell, J. F. (1997). A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environmental Conservation*, 24, 38–49.
- Foster, D. R., & Aber, J. D. (2004). *Forest in time: Environmental consequences of 1,000 years of change in New England*. New Haven: Yale University Press.
- Frenkel, A., & Ashkenazi, M. (2008). The integrated sprawl index: measuring the urban landscape in Israel. *The Annals of Regional Science*, 42, 99–121.
- Gregg, E. J., & Trites, A. W. (2008). A novel presence-only validation technique for improved Steller sea lion *Eumetopias jubatus* critical habitat descriptions. *Marine Ecology Progress Series*, 365, 247–261.
- Grimm, N. B., Faeth, S. H., Golubiewski, N. E., Redman, C. L., Wu, J., Bai, X., & Briggs, J. M. (2008). Global change and the ecological of cities. *Science*, 319, 756–760.
- Homer, C. G., Dewitz, J. A., Yang, L., Jin, S., Danielson, P., Xian, G., ... Megown, K. (2015). Completion of the 2011 National Land Cover Database for the conterminous United States-Representing a decade of land cover change information. *Photogrammetric Engineering and Remote Sensing*, 81, 345–354.
- Hu, Z., & Lo, C. P. (2007). Modeling urban growth in Atlanta using logistic regression. *Computers, Environment and Urban Systems*, 31, 667–688.
- Intergovernmental Panel on Climate Change [IPCC]. (2007). Climate change 2007, Synthesis Report. 107 p. [online] URL: [http://www.ipcc.ch/publications\\_and\\_data/publications\\_ipcc\\_fourth\\_assessment\\_report\\_synthesis\\_report.htm](http://www.ipcc.ch/publications_and_data/publications_ipcc_fourth_assessment_report_synthesis_report.htm) (Date accessed: 25.03.11).
- Irani, F.M., & Claggett, P. (2010). Chesapeake Bay Watershed Land Cover Data. Series. US Geological Survey Data Series, 505.
- Jantz, C. A., Goetz, S. J., & Shelley, M. K. (2003). Using the SLEUTH urban growth model to simulate the impacts of future policy scenarios on urban land use in the Baltimore-Washington metropolitan area. *Environment and Planning B: Planning and Design*, 2003(30), 251–271.
- Johnson, C. J., Nielsen, S. E., Merrill, E. H., McDonald, T. L., & Boyce, M. S. (2006). Resource selection functions based on use-availability data: theoretic motivation and evaluation methods. *Journal of Wildlife Management*, 70, 347–357.
- Jones, K. R., Watson, J. E. M., Possingham, H. P., & Klein, C. J. (2016). Incorporating climate change into spatial conservation prioritization: A review. *Biological Conservation*, 194, 121–130.
- Li, X., & Gong, P. (2016). Urban growth models: progress and perspective. *Science Bulletin*, 1–16.
- Lin, L. I. (1989). A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, 1, 255–268.
- Matthews, R. B., Gilbert, N. G., Roach, A., Polhill, J. G., & Gotts, H. M. (2007). Agent-based land-use models: a review of applications. *Landscape Ecology*, 22, 1447–1459.
- McGarigal, K., Compton, B. W., Plunkett, E. B., Deluca, W. V., & Grand, J. (2017). Designing sustainable landscapes: project overview. Report to the North Atlantic Conservation Cooperative, US Fish and Wildlife Service, Northeast Region. [online] URL: [https://scholarworks.umass.edu/designing\\_sustainable\\_landscapes/](https://scholarworks.umass.edu/designing_sustainable_landscapes/).
- McGarigal, K., Compton, B. W., Plunkett, E. B., Deluca, W. V., Grand, J., Ene, E., & Jackson, S. D. (2018). A landscape index of ecological integrity to inform landscape conservation. *Landscape Ecology* (submitted for publication).
- Meentemeyer, R. K., Tang, W., Dornig, M. A., Vogler, J. B., Cunniffe, N. J., & Shoemaker, D. A. (2013). FUTURES: multilevel simulations of emerging urban-rural landscape structure using a stochastic patch-growing algorithm. *Annals of the Association of American Geographers*, 103, 785–807.
- Millard, M. J., Czarnecki, C. A., Morton, J. M., Brandt, L. A., Briggs, J. S., Shipley, F. S., ... Taylor, J. (2012). A national geographic framework for guiding conservation on a landscape scale. *Journal of Fish and Wildlife Management*, 3, 175–183.
- Nakicenovic, N., Alcamo, J., & Davis, G. (2000). *Special report on emissions scenarios: A special report of working group III of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press 599 p. [online] URL: <http://www.grida.no/climate/ipcc/emission/index.htm> (Date accessed: 25.03.11).
- Newbold, T., Hudson, L. N., Hill, S. L. L., Contu, S., Lysenko, I., Senior, R. A., ... Purvis, A. (2015). Global effects of land use on local terrestrial biodiversity. *Nature*, 520, 45–50.
- Parker, D. C., Manson, S. M., Janssen, M. A., Hoffmann, M. J., & Deadman, P. (2003). Multi-agent systems for the simulation of land-use and land-cover change: a review. *Annals of the Association of American Geographers*, 93, 314–337.
- R Core Team. (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: <http://www.R-project.org/>.
- Reis, J. P., Silva, E. A., & Pinho, P. (2015). Spatial metrics to study urban patterns in growing and shrinking cities. *Urban Geography*, 37, 246–271.
- Santé, I., García, A. M., Miranda, D., & Crecente, R. (2010). Cellular automata models for the simulation of real-world urban processes: a review and analysis. *Landscape and Urban Planning*, 96, 108–122.
- Seto, K. C., Guneralp, B., & Hutyra, L. R. (2012). Global forecasts of urban expansion to 2030 and direct impacts on biodiversity and carbon pools. *Proceedings of the National Academy of Sciences USA*, 109, 16083–16088.
- Sim, S., & Mesev, V. (2014). Measuring and modeling of urban growth and its impacts on vegetation and species habitats in greater Orlando, Florida. *International Journal of Geospatial and Environmental Research*, 1(1) Article 5.
- Sohl, T. L., Saylor, K. L., Drummon, M. A., & Loveland, T. R. (2007). The FORE-SCE model: a practical approach for projecting land cover change using scenario-based modeling. *Journal of Land Use Science*, 2, 103–126.
- Sohl, T., & Saylor, K. (2008). Using the FORE-SCE model to project land-cover change in the southeastern United States. USGS Staff – Published Research. Paper 572. [online] URL: <http://digitalcommons.unl.edu/usgsstaffpub/572> (Date accessed: 1.03.17).
- Sudhira, H. S., Ramachandra, T. V., & Jagadish, K. S. (2004). Urban sprawl: metrics, dynamics and modeling using GIS. *International Journal of Applied Earth Observation and Geoinformation*, 5, 29–39.
- The Nature Conservancy (TNC) secured lands. [online] URL: <https://www.conservationgateway.org/ConservationByGeography/NorthAmerica/UnitedStates/edc/reportsdata/terrestrial/secured/Pages/default.aspx> (Date accessed: 25.03.11).
- Theobald, D. (2005). Landscape patterns of exurban growth in the USA from 1980 to 2020. *Ecology and Society*, 10(32) URL: <http://www.ecologyandsociety.org/vol10/iss1/art32/> (Date accessed: March 1, 2017).
- Triantakostas, D., & Mountrakis, G. (2012). Urban growth prediction: a review of computational models and human perceptions. *Journal of Geographic Information System*, 4, 555–587.
- Veldkamp, A., & Fresco, L. O. (1996). CLUE-CR: an integrated multi-scale model to simulate land use change scenarios in Costa Rica. *Ecological Modeling*, 91, 231–248.
- Verburg, P. H., DeKoning, G. H. J., Kok, K., Veldkamp, A., & Bouma, J. (1999). A spatial explicit allocation procedure for modeling the pattern of land use change based upon actual land use. *Ecological Modeling*, 116, 45–61.
- Verburg, P. H., Schot, P. P., Dijst, M. J., & Veldkamp, A. (2004). Land use change modeling: current practice and research priorities. *GeoJournal*, 61, 309–324.
- Verburg, P. H., Soepboer, W., Limpiada, R., Espaldon, M. V. O., Sarifa, M., & Veldkamp, A. (2002). Land-use change modeling at the regional scale: the CLUE-S model. *Environmental Management*, 30, 391–405.
- Wear, D.N. (2011). Forecasts of county-level land uses under three future scenarios: a technical document supporting the Forest Service 2010 RPA Assessment. Gen. Tech. Rep. SRS-141. Asheville, NC: U.S. Department of Agriculture Forest Service, Southern Research Station. 41 p.
- Wear, D. N., Pye, J., & Ritters, K. (2004). Defining conservation priorities using fragmentation forecasts. *Ecology and Society*, 9(4) [online] URL: <http://www.ecologyandsociety.org/vol9/iss5/art4/> (Date accessed: March 1, 2017).
- Wickham, J. D., Stehman, S. V., Gass, L., Dewitz, J., Sorenson, D. G., Granneman, B. J., Poss, R. V., & Baer, L. A. (2017). Thematic accuracy assessment of the 2011 National Land Cover Database (NLCD). *Remote Sensing of Environment*, 191, 328–341.
- White, R., & Engelen, G. (2000). High-resolution integrated modeling of the spatial dynamics of urban and regional systems. *Computers, Environment and Urban Systems*, 24, 383–400.

Wu, F. (2002). Calibration of stochastic cellular automata: the application to rural–urban land conversions. *International Journal of Geographic Information Science*, 16, 795–818.

Xiang, W.-N., & Clarke, K. C. (2003). The use of scenarios in land-use planning. *Environment and Planning B: Planning and Design*, 30, 885–909.

Zhao, Y., & Murayama, Y. (2007). A new method to model neighborhood interaction in cellular automata-based urban geosimulation. In: Y. Shi et al. (Eds.). ICCS 2007, Part II, LNCS 4488, pp. 550–557.