Designing Sustainable Landscapes: Modeling Urban Growth

A project of the University of Massachusetts Landscape Ecology Lab

Principals:

- Kevin Mcgarigal, Professor
- Brad Compton, Research Associate
- Ethan Plunkett, Research Associate
- Bill DeLuca, Research Associate
- Joanna Grand, Research Associate

With support from:

- North Atlantic Landscape Conservation Cooperative (US Fish and Wildlife Service, Northeast Region)
- Northeast Climate Science Center (USGS)
- University of Massachusetts, Amherst

Report date: 20 April 2018

Reference:

McGarigal K, Plunkett EB, Willey L, Compton BW, DeLuca WV, and Grand J. 2017. Designing sustainable landscapes: modeling urban growth. Report to the North Atlantic Conservation Cooperative, US Fish and Wildlife Service, Northeast Region.

Table of Contents

1 Problem Statement

One of the critical landscape change drivers in the Northeast (and elsewhere) is human development associated with urban growth. Indeed, urban growth is perhaps the principal driver of permanent landscape change affecting both the integrity of ecological systems and the capacity of the landscape to support species. Thus, any assessment of the current and/or future landscape to support biodiversity must account for urban growth. In this document, we describe the approach for modeling urban growth in Landscape Change, Assessment and Design (LCAD) model of the Designing Sustainable Landscapes (DSL) project (McGarigal et al 2017).

2 Solution Statement

We sought an urban growth model that would meet the following objectives:

- Ensure the model can be applied over the entire Northeast;

- Allow multiple types of land-use change, including the development of new low-, medium-, and high-intensity land-use and the transition of low to medium, low to high, and medium to high;

- Allow for the implementation of multiple urban growth scenarios;

- Model both the pattern of growth and amount of growth;

- Enforce non-stationarity of urban growth patterns and amounts across space and time; and

- Use an empirically-based approach that makes the best use of available data.

To meet these objectives, we developed a custom urban growth model (SPRAWL) to project land-use conversion over time under different scenarios of demand for land-use conversion and sprawl intensity. The model consists of a training phase and an application phase with four inter-dependent model components through which empirical models trained by historical growth information are applied in a non-stationary way across the Northeast using a unique matching approach.

The *training phase* utilizes historical data from three training regions, located within the Northeast, for which spatially explicit, decade long, land-use data are available. The four components of the *application phase* were developed independently and then combined to form the integrated model. The multi-part structure of this model (with separate demand, suitability and allocation components) is loosely based on the CLUE-s model framework (Verburg et al. 2002), which has since been modified and utilized by other land-use/land-cover change models, including the USGS FORE-SCE model (forecasting scenarios of land cover change)(Sohl et al. 2007, Sohl and Sayler 2008). In an effort to install non-stationarity across space and time, and because we are implementing this strategy throughout the entire Northeast, we adopted an additional matching element that is unique to our approach, details of which are outlined below.

3 Key Features

3.1 Training phase

The SPRAWL model utilizes historical training data to characterize urban growth amount and spatial patterns for six different transition types (e.g., undeveloped to low-intensity development, medium-intensity to high-intensity development). The training data were taken from a subset of the Northeast, specifically Maine and Massachusetts, and the Chesapeake Bay (**Fig. 1**). The primary training data source for Maine and Massachusetts was the National Oceanic and Atmospheric Administration (NOAA) Coastal Services Center Coastal Change Analysis Program (C-CAP) data which was available 1996 and 2006 (http://www.csc.noaa.gov/crs/lca/northeast.html). We also utilized Chesapeake Bay Watershed Landcover Data Series (CBLCD) available from 1984 and 2006 (ftp://ftp.chesapeakebay.net/Gis/CBLCD_Series/), which was based partly on the C-CAP data as well.

We divided each of the three training regions into non-overlapping training windows ~15 km on a side (**Fig 1**). Each training window occupied a position in a standardized two-dimensional state space defined by the intensity of development and open water. We located 16 uniformly distributed "model points" or locations in this state space (**Fig. 2**) and fit separate binary logistic regression models for each of the six transition types to a set of training points (i.e., cells of the corresponding transition type at least 150 m apart matched with an equal number of randomly selected "available" cells) from the training windows located near each model point in the model state space. The predictor variables in the logistic regression models included a variety of measures of intensity of open water, roads and development at different scales, as well as a transformation of slope and distance to nearest road.
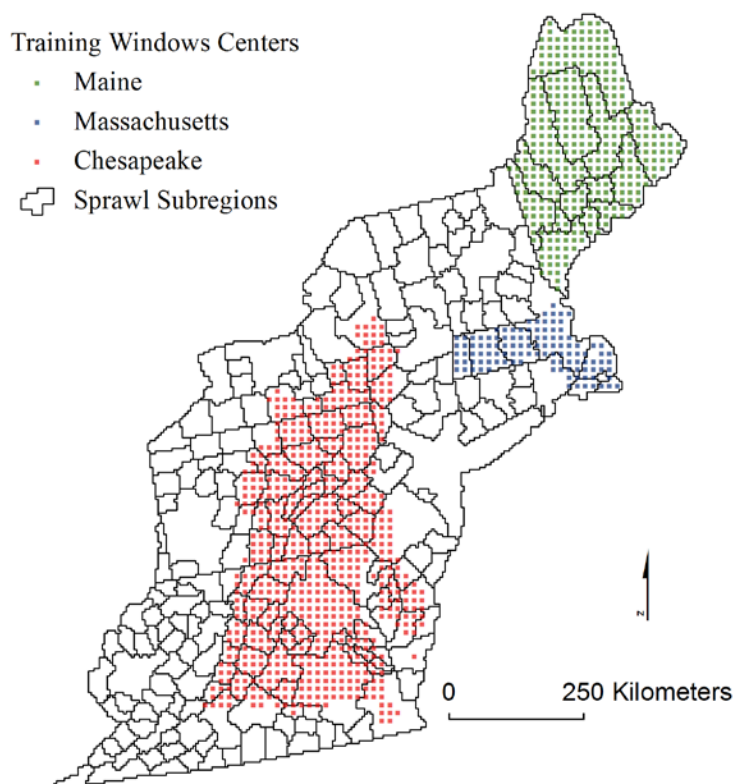

Figure 1. Training data for the urban growth model were derived from the three areas depicted within the Northeast. Solid line depict our gridded (~5 km resolution) version of Core Base Statistical Areas (CBSAs) and counties (i.e., subregions) across the Northeast for which we designate the demand (in cells) for urban growth based on the 2010 RPA forecasts (Wear 2011).
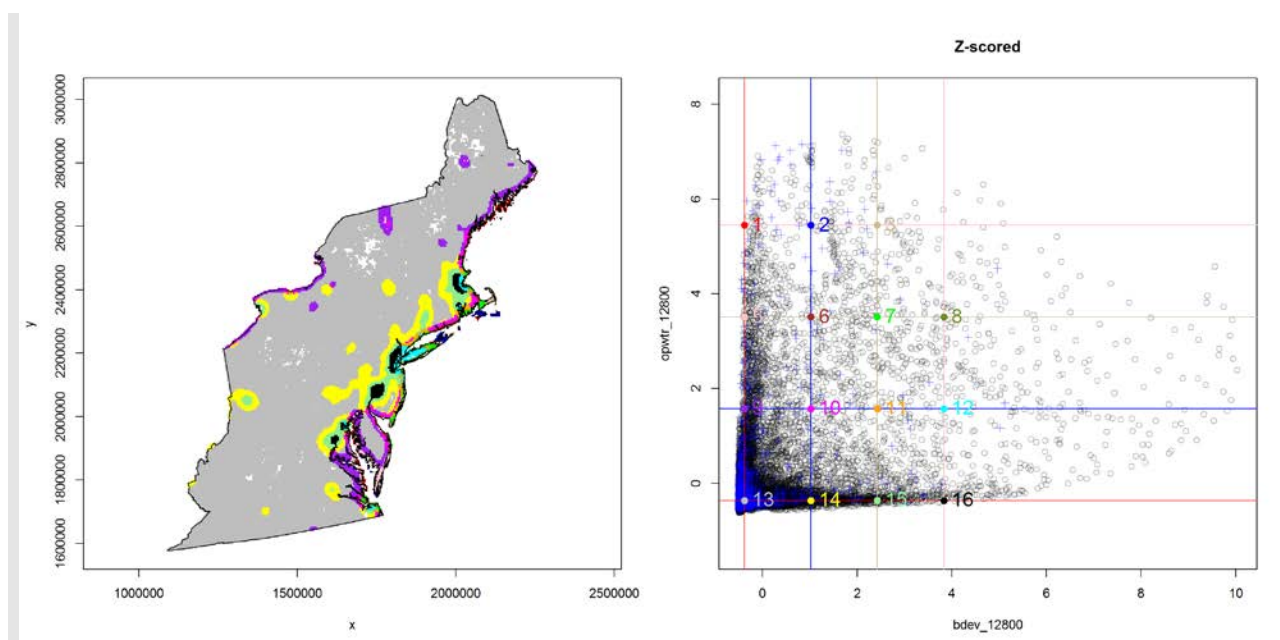
Figure 2. Right figure depicts the two-dimensional model state-space for predicting growth patterns; it is based on Gaussian kernel (bandwidth=12.8 km) intensity of development (x-axis) and open water (y-axis) within which training windows (~15 km on a side) are located across three training areas in the Northeast (**Fig. 1**). Left figure shows the spatial distribution (at the ~5 km application pane resolution) of the 16 model points depicted in the model state space.

Ultimately, we ended up with a logistic regression model to predict each of the six transition types for each of 12-16 model points, depending on transition type, uniformly distributed throughout the state space defined by the intensity of development and open water (**Fig. 3**).

In addition, for each training window we also computed the historical distribution of transition types (i.e., the proportion of total transitions comprised of each of the six transition types), the distribution of observed sizes of disjunct development patches, and the total amount (in cells) of historical development - the 'match amount', for use in the simulation (see below).

3.2 Application phase

The SPRAWL model consists of several interacting components. The basis for the 'current', or initial, land-use condition in the LCAD simulation (set to be roughly the year 2010) is the set of developed landcover classes in DSLland (see DSLland document, McGarigal et al 2017), including low-, medium- and high-intensity development derived from the 2011 National Landcover Dataset (NLCD).

For purposes of the SPRAWL model, we subdivide the entire Northeast region into non-overlapping square application "panes" ~5 km 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). Given this spatial template, the urban growth model is implemented as follows:

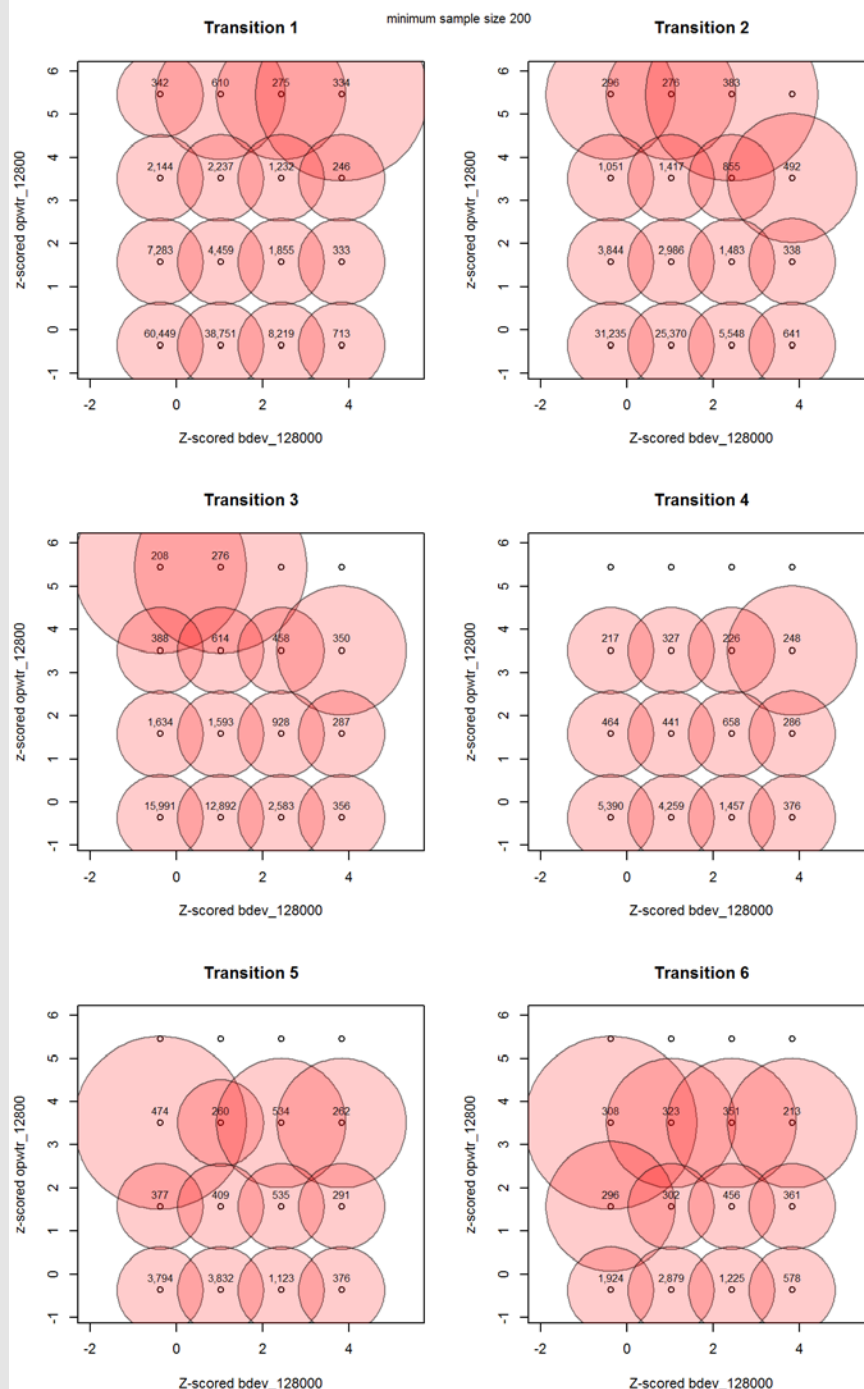**Figure 3.** The radius (1, 1.5, or 2 standard deviations) around each of the 16 model points in the two-dimensonal state space, as defined in **figure 2**, and the number of sample points included in each logistic regression model for each of the six development transitions (see text for details). Missing circles represent model points for which the minimum sample size of 200 was not met within 2 standard deviations.

1) *Demand* — To begin, we establish the "demand" for additional cells of urban land-uses (including low-, medium-, or high-intensity development) at each 10-year timestep from 2010 to 2080 for each application pane based on downscaled county-level forecasts derived for a U.S. Forest Service 2010 Resources Planning Act (RPA) assessment (Wear 2011).

2) *Matching* — Next, to allocate the total demand (in cells) within each census Core Base Statistical Areas (CBSA) or county to each application pane at each timestep, each application window is matched to the three most similar training windows based on geographic proximity and four landscape metrics based on the intensity of development, roads, and open water, and the density of roads within the window.

3) *Allocation* — Once each application window is matched to three training windows, we allocate the total demand (in cells) within the corresponding CBSA or county for the current timestep to each application pane based on how much of the historical development was allocated to the training windows. The end result is that the total demand (in cells) for each CBSA or county 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. Lastly, the demand (in cells) in each application pane for the current timestep is allocated among the six transition types based on the historical distribution in the matched training windows.

4) *Suitability* — 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 two-dimensional model state-space described above. We use these weighted-average models to compute the relative probability (i.e., "suitability") of each transition type for each 30 m cell in the application pane.

Given the demand (in cells) for each transition type allocated to each application pane for the current timestep and the corresponding suitability surfaces, we simulate actual development for each transition type, as follows:

1) Randomly select a cell to initiate the disturbance based on the relative probability (i.e., suitability) surface fit for that transition;

2) randomly draw a patch size from the observed distribution of patch sizes in the three matching training windows for the corresponding transition type.

3) spread outward from the initiation cell with a resistant kernel, where resistance is based on the inverse of the corresponding probability of transition for each neighboring cell, until the randomly selected patch size is met, allowing patches to extend across the boundaries of the focal application pane.

4) repeat the process above, building development patches sequentially until the total allocation of cells for the transition type is exhausted in the application pane.

Urban growth scenarios --Urban growth scenarios can be implemented by varying the overall amount of land that is developed (i.e., increase or decrease the amount of development relative to the RPA forecasted amount) or by adjusting 'sprawl dial' that

determines how contagious (compact vs. "sprawlly") growth patterns will be at the broad scale, compared to historic trends.

In summary, the urban growth model acts as a disturbance process on the landscape, realizing development at the 30 m cell and patch level in each 5 km pane at each 10-year timestep until the allocated number of cells to be disturbed have been exhausted. The types of disturbance transitions (e.g., undeveloped to low intensity development, or medium- to high-intensity development) are allocated proportionately to that observed historically in the most similar training windows. The patterns of development for each transition type are modeled to reflect the historical patterns that occurred in landscapes having a similar landscape context. The sizes of patches developed are chosen to reflect the distribution observed historically in the most similar training windows. At the end of each 10-year timestep, once growth is realized, the resulting urban grid is fed back into the beginning of the process for the next timestep. Importantly, at each timestep each application window is matched to three new training windows and the weights assigned to each training model are recalculated . In this way, the model is non-stationary across space and time; as a window becomes more urbanized in the future, its growth patterns change to match the way more urbanized windows grew historically, but all subject to the projected demand for growth at the CBSA or county level.

4 Detailed Description of Process

The SPRAWL model consists of an extensive training phase and an application phase. A detailed description of each phase follows:

4.1 Training Phase

4.1.1 Training data

We selected three sub-regions across the Northeast to be used as training areas: Maine, Massachusetts, and the Chesapeake Bay (**Fig. 1**). The primary training data sources were as follows:

- *Maine and Massachusetts* — National Oceanic and Atmospheric Administration (NOAA) Coastal Services Center Coastal Change Analysis Program (C-CAP) data. These data were available for different time periods for many coastal areas within the U.S. However, not all inland areas were available. Much of the Northeast was available in grid format for the 1996 and 2006 timesteps: http://www.csc.noaa.gov/crs/lca/northeast.html.

- *Chesapeake Bay* — Chesapeake Bay Watershed Landcover Data Series (CBLCD) available from 1984 and 2006: ftp://ftp.chesapeakebay.net/Gis/CBLCD_Series/, 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 region. However, it is important that the training data be representative of conditions throughout the region and directly comparable to a land-use layer for the "current" (or starting) condition that is available throughout the region. The three training areas were chosen because they each have a

relatively long period of historical land-use GIS layers available (1996–2006 in Maine and Massachusetts from NOAA C-CAP, and 1984–2006 in Chesapeake Bay from CBLCD) that have been analyzed to evaluate change over multiple urban classes (low-, medium-, and high-intensity urban). Together, they represent varied amounts of current land-use intensity and historic land-use conversion that likely represent the full gradient found across the entirety of the Northeast.

In many ways, the C-CAP data series is directly comparable to the 2006 National Land Cover Database (NLCD; Fry et al. 2011), which is available at the 2006 timestep throughout the entire Northeast and is the most consistent and spatially comprehensive product of its kind currently available. In all three training regions, the C-CAP data are available as 30 m resolution raster GIS grids classified into various land cover categories. Like the NLCD, the C-CAP developed land-use category is broken into four classes: high, medium, and low urban intensity, and developed open space.

We examined each of the developed classes in comparison with aerial imagery for both NLCD and C-CAP, and determined that the developed open space class was unreliable and inconsistently characterized, and it had the lowest producer's accuracy of any of the developed categories in the NLCD (Wickham et al. 2010), so we treat it as static in the urban growth model; i.e., the footprint of land classified as developed open space will remain constant during an urban growth simulation. The other three developed land classes generally perform better (Wickham et al. 2010), are fairly consistent between C-CAP and NLCD, and crosswalk directly.

4.1.2 Development transitions

For each training region, we compared the grid from the earlier timestep (e.g., the 1996 C-CCAP grid) to the grid from the later timestep (e.g., the 2006 C-CAP grid), and a new *change grid* was developed with each cell defined as unchanged or as one of the following six transition types:

> *transition 1* = undeveloped to low-intensity development;
> *transition 2* = undeveloped to medium-intensity development;
> *transition 3* = undeveloped to high-intensity development;
> *transition 4* = low-intensity development to medium-intensity development;
> *transition 5* = low-intensity development to high-intensity development; and
> *transition 6* = medium-intensity development to high-intensity development.

As mentioned above, the developed open space land-use class was inconsistent and therefore treated as static in the model, so it is not considered here. In addition, note that we did not consider transitions to lower urban intensities or transitions back to a non-urban or undeveloped condition. While these transitions do occur, we believe them to be extremely rare (and thus of minor importance in the urban growth model) and, moreover, unreliably assessed from the training data.

4.1.3 Training windows

We divided each training region into non-overlapping square training "windows" approximately 15km on a side (498 30 m cells, or 14.94 km on a side)(**Fig. 1**). 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 influence urban growth and keeping the windows large enough to avoid the idiosyncrasies of very small landscapes and to reduce computational time.

For each training window, we used the change grid to calculate the historical annual urban growth rate. Specifically, we calculated and stored for later use in the Matching algorithm (see below) three statistics for each training window:

1) *Rate of new urban development* — we calculated this by counting the number of cells that underwent new development during the training period (i.e., the total number of 30 m cells in the window that underwent transition types 1, 2, or 3), divided by the length of the training period.

2) *Distribution of transition types* — we calculated the proportion of the total number of cells exhibiting any transition (transition types 1-6) during the training period comprised of each of the six transition types.

3) *Patch size distribution* — we also tabulated the distribution of development patch sizes, where a patch was defined as a set of contiguous cells (based on the eight neighbor rule) that underwent the same development transition.

4.1.4 Suitability models

For each training window, we computed the Gaussian kernel (12.8 km bandwidth) intensity of development and open water, and converted these values to z-scores (i.e., mean=0 and standard deviation=1). Thus, each training window occupied a position in a standardized two-dimensional state space defined by the intensity of development and open water (**Fig. 2**). We located 16 uniformly distributed "model points" or locations in this state space. For each of these model 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., cells of the corresponding transition type at least 150 m apart to avoid pseudo-replication, matched with an equal number of randomly selected "available" cells) from the training windows located within 1, 1.5 or 2 standard deviations from the model point in the model state space. We used the smallest of the three standard deviations needed to achieve a minimum of 200 training points (half of which experienced that transition and half of which were available points) in order to have the fitted model most closely reflect landscape conditions described by that model point in the two-dimensional state space. 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.

Predictor variables in the logistic regression models included:

- Gaussian kernel (bandwidths=100, 800 and 3,200 m) intensity of open water;
- Gaussian kernel (bandwidth=800 m) intensity of primary and secondary roads;
- Gaussian kernel (bandwidth=800 m) intensity of all roads except motorways;
- Gaussian kernel (bandwidth=3,200 m) intensity of all roads except motorways;
- Gaussian kernel (bandwidths=400 and 3,200 m) intensity of weighted development (weights for low-intensity development=1, medium-intensity development=2, high-intensity development=3; NA on all cells not eligible for development);
- Transformed slope based on a univariate logistic regression model (**Fig 4**); and

- Transformed distance to the nearest road (excluding motorways) based on a univariate logistic regression model (**Fig. 5**).

Note, the selection of the final suite of predictor variables listed above was informed by an earlier (phase 1 of the DSL project) comprehensive and objective hierarchical model selection process involving 103 different predictors representing combinations of a large number of candidate variables and wide range of spatial scales.

In addition, the transformations applied to slope and distance to road were based on preliminary evaluations of results using the untransformed variables. In particular, even with a large effect size (i.e., magnitude of the standardized regression coefficient) for untransformed distance to roads, due to the disproportionate amount of land away from roads, we ended up simulating too much new development away from roads (by comparison to the distribution observed in the training data). Similarly, we were simulating too much development on steep slopes. In both cases, these transformations ensured that new simulated development better reflected the historical patterns with respect to distance from road and slope.

Despite using the transformed distance to the nearest road in the logistic regression models, we noted that our logistic regression models were not forcing enough of the new development (transitions 1-3) to be close to roads. This was due to bias in the training data resulting from


Figure 4. Gussian transformation of slope (%) for use in logistic regression models to predict probability of urban development.


Figure 5. Gussian transformation of distance (m) to road (excluding motorways) for use in logistic regression models to predict probability of urban development.

geo-processing to eliminate the confounding of roads and development in the C-CAP and CBLCD data. To adjust for this bias, we multiplied the logistic functions for transitions 1-3 by a negative exponential function of distance to nearest road (exponent = -0.011965238). The latter was fit to the distribution of new developments (transitions 1-3 pooled) observed in the hindcast dataset we developed to validate the urban growth model (see below). We did not apply this road adjustment to transitions 4-6.

Ultimately, we ended up with a logistic regression model (a.k.a. suitability model) to predict each of the six transition types for each of 12-16 model points (depending on transition type) uniformly distributed throughout the state space defined by the intensity of development and open water (**Fig. 3, Appendix A**). During the application phase (see below), these models are used in the Suitability algorithm to determine the pattern of growth for each transition type in each application window.

Figure 6. Application "panes" (~5 km on a side) embedded within overlapping application "windows" (~15 km on a side) as used in the Application phase of the urban growth model.

4.2 Application Phase

The application phase involves simulating urban growth according to a user-specified scenario based on an algorithm consisting of several interacting components, as described below. The basis for the "current", or initial, land-use condition in the urban growth model (set to be roughly the year 2010) is the set of developed landcover classes in DSLland, including low-, medium- and high-intensity development derived from the 2011 National Landcover Dataset (NLCD).

For purposes of simulating urban growth, we subdivided the entire Northeast region into non-overlapping square application "panes" ~5 km (166 cells, or 4,980 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, similar to the training windows)(**Fig. 6**).

4.2.1 Determining the demand

To begin, we established the "demand" for additional cells of urban land-uses (including low-, medium-, or high-intensity development) at each 10-year timestep from 2010 to 2080. The demand dictates the overall *amount* (in cells) of urban land-uses to allocate throughout the area of interest in each timestep.
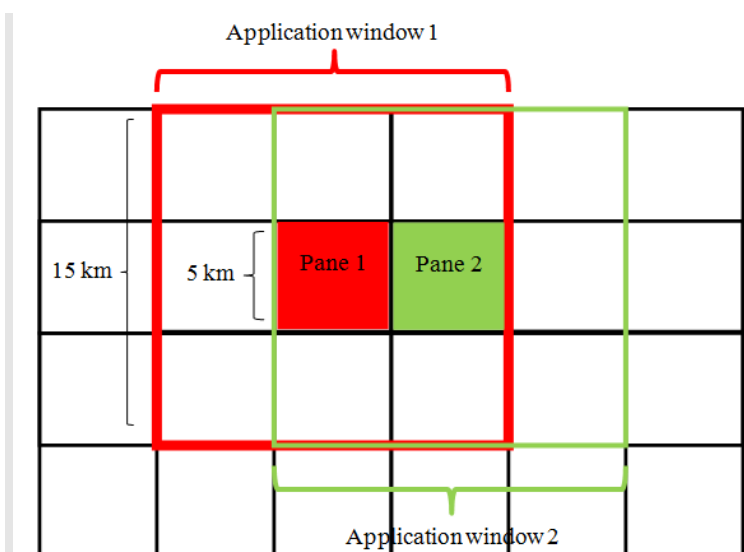
We based the demand for urban growth at each future timestep on county-level forecasts derived for a U.S. Forest Service 2010 Resources Planning Act (RPA) assessment (Wear 2011), which can be found on the last sheet of this file:

https://www.fs.fed.us/research/docs/rpa/2010/2010RPA_Assessment_landuse projections_county_final.xlsx

For our purposes, it was necessary to resolve three issues with the RPA data:

1) RPA projections were for 424 counties in the Northeast, but some of these were merged counties. There was no data on the counties that were merged out of existence and it was not 100% clear for all of these counties which of the neighboring counties they were merged with. With most counties it was obvious; for example, in Virginia some major cities are essentially separate counties with their own FIPS code (county level ID). These cities are distinct from but completely surrounded by a county and so clearly would have been merged into that county. For other counties that were merged out of existence in the RPA dataset it was harder to tell which of the adjacent counties they were merged with. Ultimately, based on the geography we manually generated a crosswalk between the full set of counties and the RPA reduced set containing some merged counties.

2) Twenty-four counties in the RPA assessment appeared to have missing data for all timesteps (for unknown reasons); the corresponding rows in the tables indicated zero development. We interpreted these projections to be erroneous and replaced these zeros with imputed values based on the mean of the 10 counties that were most similar in population density.

3) RPA projections extended only to the year 2060. We extrapolated the values reported from 2050-2060 to the expected development in 2070 and 2080.

4) RPA projections were made for three scenarios are linked to globally consistent and well-documented scenarios used in the Intergovernmental Panel on Climate Change (IPCC) 4th Assessment (AR4) (IPCC 2007). The scenarios were the A1B, A2, and B2 scenarios, which include a range of future global and U.S. socioeconomic and climate conditions likely to affect future U.S. resource conditions and trends (Nakicenovic and others 2000). We decided to use the average projections across these three scenarios to better reflect the uncertainty in future conditions.

With these issues resolved, 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 (**Fig. 1**). 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 CBSA or county in our gridded version by multiplying the forecasted rate of development by the count of land cells within the gridded CBSA or county. 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.

4.2.2 Simulating urban growth

Based on the total demand (above), we simulate urban growth for each timestep as follows:

1) *Matching* — First, each application window is matched to the three most similar training windows (from any of the three training regions) based on geographic proximity and four landscape metrics:

   - Gaussian kernel (bandwidth=12,800 m) intensity of development;
   - Gaussian kernel (bandwidth=12,800 m) intensity of all roads, except motorways;
   - Gaussian kernel (bandwidth=12,800 m) intensity of open water; and
   - Density of roads (except motorways) within the window.

   Note, the selection of the landscape structure variables listed above was informed by an earlier (phase 1 of the DSL project) comprehensive and objective model selection process involving 50 different variables to determine which combination of variables had the greatest ability to predict the amount and pattern of future growth in each training window.

   In addition, 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 prevent major boundary effects across application panes. In this way, each application pane shares 2/3 of its window with its four nearest neighboring panes (**Fig. 6**), preventing abrupt differences from occurring between neighboring panes.

   For each application window and training window, we compute the four landscape metrics listed above and convert the values to z-scores (i.e., 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 (for reasons discussed below). 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 (**Fig. 7**). The resulting adjusted similarity scores are rank ordered and the three most similar (weighted by proximity) training windows are matched for each application window.

   Note, due to the lack of regionally available and consistent spatial data, many of the local socio-economic drivers (e.g., zoning laws, median income, etc.) that influence land-use change have not been incorporated into the model. We recognize that application windows are likely to grow more like windows near themselves simply because they have the same 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.

2) *Allocation* — Once each application window is matched to three training windows, we allocate the total demand (in cells) for the current timestep within each gridded CBSA or county to each application pane. To do this, we calculate the average total amount of historical development (the "match amount") in the three training windows. This match amount is subsequently adjusted to reflect the proportion of the application window that is buildable and the proportion of the buildable in the window that is in the central pane. Specifically, the match amount is multiplied by the proportion of the window that is buildable (i.e., not water, wetland, secured, or already developed) and the proportion of the buildable in the window that is in the central pane. We also adjust this match amount as necessary to ensure that no

Figure 7. Adjustment multiplier for similarity scores between application windows and training windows in which the closest training window receives a delta=0 and no adjustment, and as the delta (difference from the closest window) increases the adjustment multiplier decreases according to the logistic Gaussian function shown.

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 development threshold was based on the 99th 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 timestep is computed by dividing the pane's interim match amount by the total interim match amount across all panes in the corresponding CBSA or county. In this manner, the total demand (in cells) for each CBSA or county 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 timestep 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
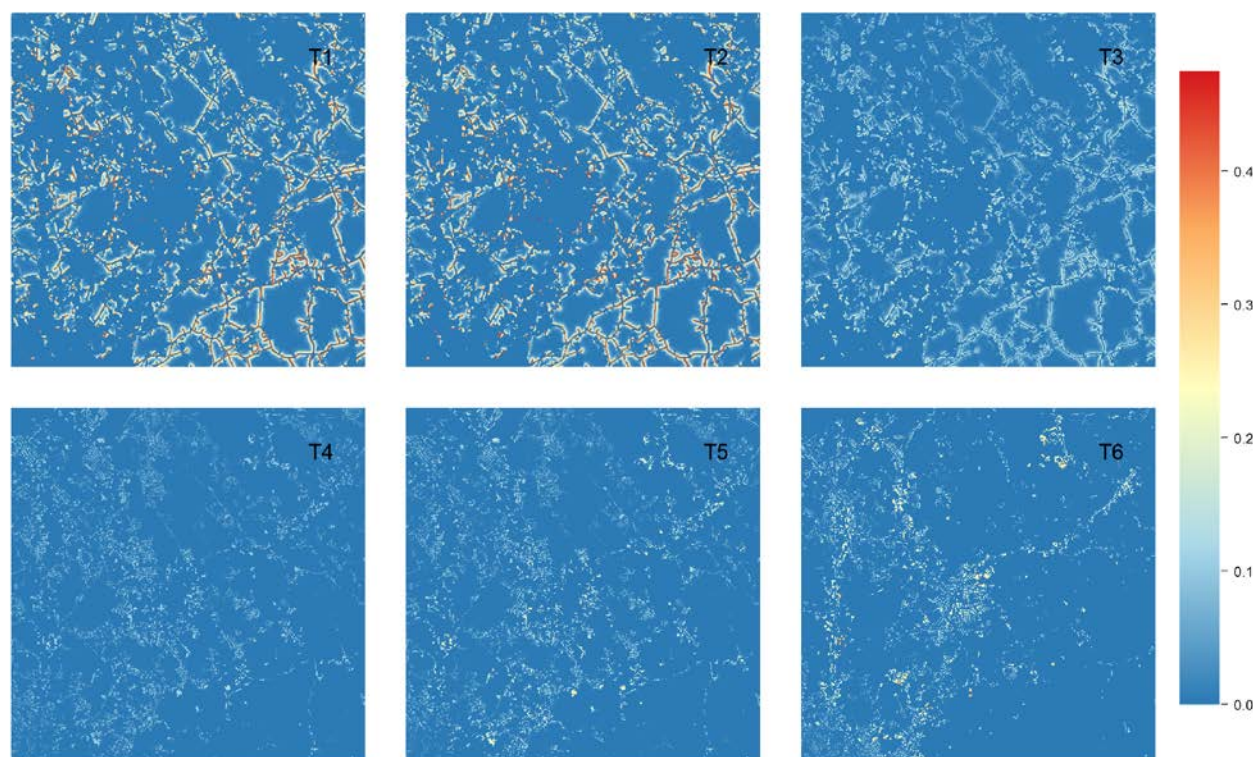
Figure 8. Probability of development (i.e., suitability) at the 30 m cell level for a random application window and timestep for each of six transition types (see text for details) for a random location in the Northeast.

allocation to the pane, and the ratio among all six transitions made to match the historical ratios in the matched training windows.

3) *Suitability* — Once we determine the number of cells needed for each transition type in each application pane for the current timestep, we determine the probability of each transition type occurring at the 30 m 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 two-dimensional model state-space developed during the Training phase described above. 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 30 m cell in the application pane (**Fig. 8**). Lastly, for transition types 1-3, we assign a zero value to all non-buildable cells, which includes open water, wetland, secured lands, roads and developed. 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).

Note, because the Gaussian kernel (12.8 km bandwidth) intensity of 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 time and become more like the patterns characteristic of increasingly urbanized windows. This feature ensures non-stationarity in the patterns of urban growth over time, which we deem an important and somewhat novel feature of our urban growth model.

4) *Disturbance patches* — Given the demand (in cells) for each transition type allocated to each application pane for the current timestep and the corresponding suitability surfaces, we simulate actual development for each transition type, as follows:

   a) 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, we also limit the selection of cells to those in the top 15th percentile of suitability within the application pane; thus, cells with the highest suitability values are more likely to initiate a disturbance;

   b) randomly draw a patch size from the observed distribution of patch sizes in the three matching training windows for the corresponding transition type;

   c) spread outward from the initiation cell with a resistant kernel (see technical document on integrity, McGarigal et al 2017, for a detailed description of resistant kernels), where resistance is based on the adjusted complement of the corresponding probability of transition for each neighboring cell, until the randomly selected patch size is approximately met, and allowing patches to extend across the boundaries of the focal application pane. Specifically, resistance is calculated as:

$$z_{ij} = (1 - p_{ij}) \times m$$

$$r_{ij} = z_{ij} - \min(z_{ij}) + 1$$

   where p_{ij} = probability of the j^{th} transition type for the i^{th} cell, m = resistance factor (currently set to 30), which determines the theoretical upper bound of resistance, z_{ij} = interim value, and r_{ij} resistance of the j^{th} transition type for the i^{th} cell. Thus, resistance values range from 1 to an upper bound of 31 ($m + 1$).

   Note, non-buildable cells (open water, wetland, secured lands, road and developed for transition types 1-3cells, and anything not of the focal development class for transition types 4-6) are assigned double resistance, thus the spread can cross these but with much higher resistance. We then zero out the spread kernel for non-buildable cells and determine the threshold in the spread kernel height that yields the size closest to the target patch size. All the cells in the disturbed patch over the threshold value are set to the new land-use class of the development, otherwise they retain their original value, and once a cell has transitioned in a given timestep, it is masked and not allowed to transition again during that timestep. Note, because a single resistant kernel can in some cases spread through open water, wetland, secured, road and developed, it is possible for a single disturbance to create more than one discrete development patch.
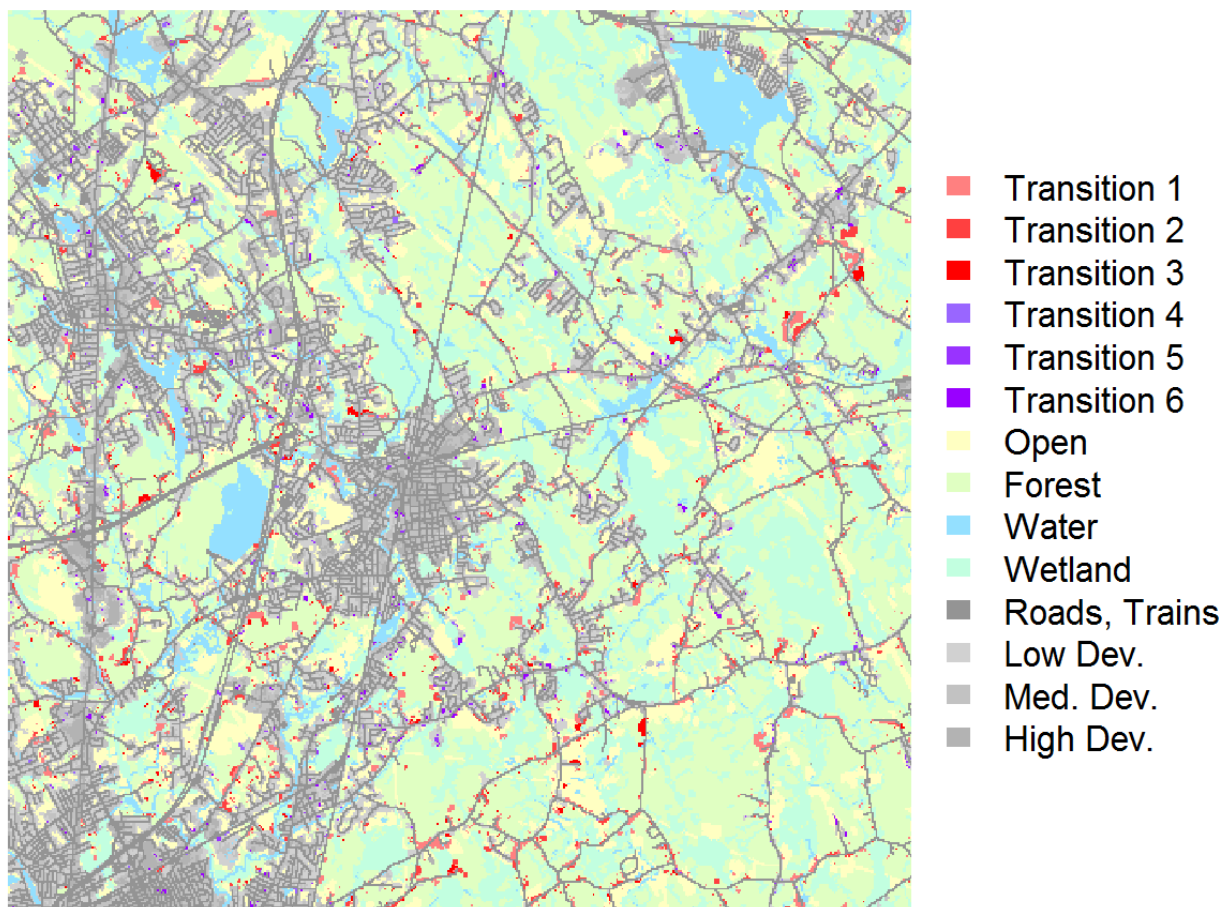
Figure 9. Simulated development transition patches for a random application window and timestep.

    d) 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. 9**).

4.2.3 Urban growth scenarios

Urban growth scenarios can be implemented in two ways:

1) *Development rate* — the overall amount of land that is developed (or the corresponding rate of development) can be modified from the baseline demand computed above; i.e., to increase or decrease the amount of development relative to the RPA forecasted amount. Note, in this scenario analysis the relative allocation of development among application panes and patterns of development at the 30 m cell level remain the same as the baseline (i.e., they reflect historical allocation and

patterns in similar landscape contexts), but the total amount of development varies (**Fig. 10**).

2) *Sprawl* — the other factor that can be adjusted is the "sprawl dial" which allows us to create scenarios that are more "sprawlly" than the historical training data. The baseline scenario acts as the status quo (sprawl level= 0) and distributes the number of new urban cells in accordance with historical patterns using the allocation process described above. For other settings of the sprawl dial a smoothed allocation is calculated and averaged with the baseline allocation. Briefly, the smoothed allocation is created with a variable bandwidth Gaussian smooth of the baseline allocation. The development allocated to each application pane in the baseline allocation is reallocated to itself and nearby panes based on a Gaussian kernel where the bandwidth ranges from 5 to 15 km, with larger bandwidths chosen when the focal pane has more existing development. Finally, a weighted average of the standard baseline allocation and the smoothed allocation is taken with the weight determined by the setting of the sprawl dial. Thus, as the sprawl dial is increased from zero, more of the allocation to urban areas is dispersed to the outlying less urban areas, which determines how contagious (compact vs. "sprawlly") growth patterns will be at the broad scale compared to historical trends (**Fig. 10**).

5 Model evaluation

It is impractical to truly validate the SPRAWL model because the outcome depends strongly on the forecasts of future development, which for the current application was based on the RPA forecasts which in turn 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, none of which can be verified without waiting to see if the predictions come true. However, hind-casting provides an alternative for evaluating model performance. For this purpose, we used the NLCD change grid for the period 2001-2011, in which each cell was coded as either unchanged or one of our six transition types, to create a year ~2000 version of our DSL land cover grid using the same process that we used to create the 2010 DSL land cover grid (McGarigal et al 2017) but using the 2001 NLCD as input instead.

Next, we used the DSL 2000 land cover grid as the initial condition and ran the SPRAWL model for a single 10-year timestep. For the demand, we computed the realized demand for new development (transition types 1-3) within each subregion by tallying the number of corresponding transitioned cells in the new DSL 2000-2010 change grid, but excluding any cells that fell in what we had mapped as secured or wetland (since we designated these as not buildable in the SPRAWL model). Thus, we set demand to equal exactly what we observed for the buildable portion of the landscape, effectively removing demand from the model evaluation (since our purpose was not to evaluate the RPA forecasts). Next, we allocated the demand and derived suitability surfaces as usual, with a minor variation in the latter. In the SPRAWL model it is unnecessary to normalize the suitability surfaces derived from the weighted-average logistic regression models before building disturbance patches because we build disturbance patches until the demand is exhausted within each pane. Here, however, for each transition type it was necessary to normalize the surface within

Figure 10. Urban growth simulated between 2010-2080 under four scenarios: (a) baseline, (b) increased demand (25% greater than baseline), (c) increased sprawl (sprawl level =0.75, reflecting an arbitrary, but relatively high level of sprawl in which new development is redistributed somewhat from locations close to urban centers to further away), and (d) combined increased demand and sprawl, for an area east of Scranton Pennsylvania on the PA, New York, New Jersey tri-state corner. Existing development in 2010 is shown in shades of blue (with darker shades indicating higher-intensity development), whereas new development simulated between 2010-2080 is shown in corresponding shades of red.

each pane to sum to 1 and multiply by the demand in number of cells, so that the final surface reflected both the allocation to the pane and the relative suitability of the cell.

Next, for each of the transition types involving new development (i.e., types 1-3), we censused cells of that transition type from the DSL 2000-2010 change grid across the entire application region and all available buildable cells (including those that got built) and recorded the probability of development from the final suitability surface for each cell (referred to below as the "census" dataset). We also randomly sampled cells of that transition type from the DSL 2000-2010 change grid across the entire application region, ensuring a minimum 150 m apart (as in the Training phase), and an equal number of random buildable cells (including in the pool of available cells those that got built) under the same proximity constraint from the DSL 2000 grid and similarly recorded the probability of development from the final suitability surface (referred to below as the "sample" dataset).

To evaluate the performance of the urban growth model, it is important to recognize that the data were appropriately considered observations of "presence-vs-available" rather than "presence-vs-absence", since the absence of a development transition at a cell doesn't mean that the cell wasn't suitable for a transition, only that by chance it wasn't developed. Thus, the cells that underwent a transition were simply a subset of those that could have undergone a transition given the same suitability. Consequently, conventional statistics for evaluating binary logistic regression models must be interpreted with caution. Nevertheless, we computed the area under the Receiver Operating Characteristic Curve (or AUC, e.g. Fielding and Bell 1997) to gauge the model's discriminatory ability, and we used the sample dataset described above for this purpose. In our context, AUC is equal to the probability that our model-derived probability of development will rank a randomly chosen transitioned cell higher than a randomly chosen available cell.

We also computed two statistics appropriate for evaluating models derived from "presence-vs-available" data, and for these statistics we used the census data described above. First, we computed the coefficient of concordance (Lin 1989) between observed and predicted based on the methods recommended by Johnson et al. (2006). Observed is the number of transitions in each probability of transition bin (n=100 equal-interval bins); predicted is the expected number in each bin given the proportion of available (represented by the random buildable) in each bin, times the bin's average probability of transition, times the total number of observed transitions. 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.

Second, we computed the weighted skewness statistic based on the method described by Gregr and Trites (2008). Briefly, this method calculates 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 transition points that are at located at the higher end of the probability scale, resulting in a more left-skewed distribution. However, to account for the distribution of available probability values across the entire application region, the count of transition points in each probability bin (n=100 equal-interval bins) is weighted by the proportion of available points that are in each bin, and then the skewness is calculated on this weighted distribution of counts. Here, we truncated

the upper 0.1% of the distribution by removing the upper bins due to very low sample sizes in the "available" counts which produced unreliable bin weights.

6 Alternatives Considered and Rejected

There are many alternative modeling frameworks available to model future urban growth including those utilizing economic/employment theory, transportation models, agent-based models, cellular automata, and hybrid models. The largest limiting factor to most of the widely used urban growth models are that they are designed and typically applied at finer spatial scales than the Northeast and for shorter projections than the 70-year projections used in the LCAD model. Such models are able to take advantage of a number of variables, particularly socio-economic variables, that are important to explaining and driving growth, but difficult to obtain in a spatially explicit and consistent format across the entire Northeast. In addition, many existing urban growth models are prohibitively computationally intensive, and could thus not be employed efficiently at the scale of the Northeast given our current computing environment.

Nonetheless, there are still dozens of modeling frameworks that have been applied across comparably large extents (e.g., SLEUTH, CLUE, DINAMICA, GEOMOD, etc.). For the most part, these models are regionally specific and would require substantial local modification and calibration in order to be implemented for the Northeast. In addition, they typically only model binary responses (i.e., developed or not). Despite the myriad of urban growth models available, none were able to meet our model objectives (outlined in Section 2): an empirically-based, non-stationary model across space and time with multiple classes of urban intensity that could be applied across the entire Northeast for a 70-year projection under multiple scenarios.

Upon settling on the new multi-tiered, non-stationary modeling approach, there were still many alternatives to each of the component modules.

6.1 Demand

After considering many alternatives to the demand module, we decide that since the amount of growth in the future is so uncertain and likely to change based on many factors beyond our understanding or predictive ability (including global and local economic factors), the amount of growth in the future was best implemented in the user-defined scenario framework, where it could be easily adjusted to reflect different future scenarios. We decided to use the RPA forecasts as our baseline scenario. Alternative approaches to determining demand, or the amount of growth in the area of interest, include:

- *Census-based* — The U.S. Census bureau has projected population sizes (as well as age classes) for the next century. They have created high, low, and medium population increase projections which could be used for multiple scenarios. These numbers could be downscaled regionally and locally based on historic allocation of population and growth. This downscaling would yield population projections in each area unit at each timestep. The major implementation constraint to this alternative is in translating population projections to land use demand projections, we would need to project residential density (people/acre) for each unit area as well. Although this could be done by trend extrapolation, or by a "sprawl index" (e.g., Frenkel and Ashkenazi

2008) in conjunction with the extensive literature on drivers of density and sprawl in the U.S., we felt this incorporated an additional layer of uncertainty into the model. It was also not directly tied to the SRES scenarios.

- *SRES-based* — By downscaling SRES scenarios, we could develop land-use demand projections for the Northeast. County-level population projections are available under four SRES scenarios through 2100 (EPA 2009). This project also estimated housing density at the 1-ha resolution and impervious surface cover at the 1-km resolution every 5 years through 2100. The USGS LandCarbon project is also currently implementing a nationwide downscaling of the SRES scenarios (Sleeter and Sohl 2010). Ultimately this approach had the same shortcoming as the previous approach, in that it required an extra step to convert population projection to land-use conversion projections.

- *Trend extrapolation from other sources* — USDA NRCS National Resources Inventory (USDA 2009) has calculated various land-uses by state in 1982, 87, 92, 97, 02, 07. This land-use trend could be extrapolated (and modified to account for alternative scenarios) going forward. The largest limitation to this dataset is that "developed land" is not further broken out. The USGS Land Cover Trends Project (e.g., Drummond and Loveland 2010) has also calculated land-uses by ecoregion in: 1973, 1980, 1986, 1992, 2000. These data are also available for trend extrapolation: http://landcovertrends.usgs.gov/download/overview.html and may be used in designing alternative scenarios that could be implemented in future phases.

6.2 Urban suitability

Many other empirically and/or expert-derived analytical approaches could have been utilized for the urban suitability module, including geographically-weighted logistic regression, polytomous logistic regression, linear or Poisson regression, frequency ratio, analytical hierarchy process, and neural networks. Many of these approaches (including neural networks) are computationally difficult, particularly across the large Northeast extent, and many have prohibitive assumptions. Logistic regression, our chosen method, is a well-known, easily implemented approach that has been shown to perform as well or better than many other analytical approaches for modeling urban growth. We ran early versions of this model using nonparametric classification trees instead of logistic regression. Although classification trees performed marginally better (correct classification rates were slightly higher), they did not perform well enough to warrant the additional computational expense required to run them the number of times required in the model (many times per 15km training and application window). Moreover, they did not produce a probability of development surface that deemed useful, not only for implementing disturbance patches, but also to serve as the basis for the integrated probability of development product (McGarigal et al 20170 that we used in several of the landscape design products.

6.3 Allocation

The allocation component could have been implemented a variety of other ways, including:

- *CLUE* — This approach utilizes a transition matrix, which dictates the probability of each possible transition between land-uses (which must be parameterized), and weights of the importance of each land-use (which is iteratively calibrated) to assign land-uses to each cell based on suitability until demand is reached.

- *USGS FORE-SCE* — This approach is used in the LandCarbon project (Soyl and Saylor 2008). This approach is similar to (and based on) CLUE, but it converts "patches" rather than cells by sampling the distribution (mean and standard deviation) of the patch size of conversions that occurred historically. (Note: they changed this to sample from a pre-determined patch library because it took too long to model organic growth around a focal cell until the correct patch size was reached using the highest probability neighboring cells).

- *Cellular automata (CA) approach* — This approach involves building a cellular automata model (where transition rules are governed by the cell neighborhood) and combining the results of the CA with those of the suitability model (e.g., White and Engelen 2000; Zhao and Murayama 2007).

- *EPA approach* — This approach uses county level population projections (from SRES scenarios) to allocate housing units using SERGoM (Theobald 2005), which uses historic housing density to allocate number of housing units and places them spatially using: land ownership, travel time, and previous growth.

Ultimately, our approach allowed us to easily combine and implement the empirically-based logistic regression models and matching algorithm, which was more non-stationary than many of the approaches outlined above, along with our resistant kernel approach to build disturbance patches, which is analogous to a cellular automata approach.

7 Major Implementation Constraints

The challenges involved in implementing the urban growth model across the Northeast are many, the largest being the availability of accurate current and historical land-use data. In order to create an empirically-based model, it was necessary to obtain historical land-use information. The NLCD is the only land cover product available across the entire Northeast with multiple urban density classes (e.g., high, medium, and low). However, while historical information for multiple urban classes was available for a 10-year time period (2001–2011), the changes in classification methodology between 2001 and 2011 resulted in too much uncertainty in the observed transitions. Because our urban growth model is non-stationary, however, it is not necessary to train the data across the entire Northeast. Rather, representative training data can be used and applied to the entire extent due to our unique matching algorithm. This design aspect of our model allowed us to use datasets that were long enough in time but smaller in spatial extent to train the model, allowing us to use the NOAA C-CAP and CBLCD datasets in place of NLCD. Still, as a starting point for the initial condition (timestep 0 in the model set to be roughly 2010) to project growth forward in time, we are limited to a consistently available land-use dataset, and the NLCD remains the only product available at the regional scale.

In examining the NLCD data it became clear that the developed classes included the footprint of roads, as well as often a rather large buffer around roads. In addition, certain

wetland surfaces had been classified as urban land-uses as well. To correct these and other errors, roads and wetlands were burned over the development classes in the NLCD in a multi-step cleaning process to create the DSLland grid.

As a result of the model being trained with C-CAP and applied to NLCD, the differences between the two datasets are extremely important, and the cleaning process had to be applied to the C-CAP training data as well to ensure that the values used for matching were as close as possible to those in the DSLland. In comparing the two layers, however, although the amount of development (or composition) was similar for the two datasets, the pattern was slightly different, causing the landscape configuration metrics to be very different in the two datasets. As a result, the configuration metrics were dropped from the analysis and matching between training and application windows relied solely on composition and location metrics. This is a major shortcoming of our current approach and one that can only be remedied with consistent historical land use/cover data over a sufficiently long time period and over the full extent of the Northeast. Should future versions of NLCD maintain consistent mapping rules (e.g., from 2011 to 2021), it may be possible to retrain the model on the changes observed in the NLCD data.

There are several other implementation constraints. The model does not allow for cells to undergo "rewilding" (i.e., conversion from a developed class to an undeveloped class), and it does not dynamically incorporate "urban open" land-uses. Both of these are a result of the training data we used. We had very limited confidence that the conversion events from developed to undeveloped land were real and could be reliably modeled.

The current model also does not incorporate the building of new roads. Clearly roads are an important driver of land-use conversion, and they are key in many aspects of this modeling approach, including the suitability surface and the matching algorithm. New roads are very difficult to predict with accuracy. Smaller roads are not necessarily required to accurately model land-use change, as the new urban cells themselves behave as roads, acting as a magnet for additional future development. New large roads are infrequent, unpredictable events and therefore we decided probably best handled on a case by case basis as custom scenarios. However, future versions of our urban growth model could strive to include modeling the development of new roads associated with urban growth.

Unfortunately, the current model is unable to account for local effects such as zoning bylaws and socioeconomic factors. We tried to indirectly incorporate these effects, but only in part and at a relatively coarse scale, by adding the geographic distance of the training windows to the application window into the matching algorithm. Although this has the unintended consequence of adding an element of stationarity to the model, we think that the benefits of including local variation outweigh the costs of increased stationarity. Additionally, although the local variation element incorporates elements of older zoning laws, it does not account for the adoption of new policies that could influence future growth patterns.

Lastly, our urban growth 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. Note, while the allocation of development at the CBSA or county level is not explicitly tied to the past (rather, it is based on the RPA

forecasts), the local allocation within CBSAs and counties is governed by the historical allocation to similar local landscapes. Unfortunately, it is impossible to predict how these local drivers might change in the future, and so using the past seems like the best, if not the only, feasible option.

8 Major Risks and Dependencies

As noted above, our modeling framework is dependent upon the 2011 NLCD product to represent developed land uses for the initial 2010 timestep, and the NOAA C-CAP and CBLCD datasets to represent historical growth patterns. Therefore, any errors associated with these datasets will be projected forward into future timesteps. We are grateful to MRLC, NOAA, USGS, and the Chesapeake Bay Program for making these datasets available.

9 Literature Cited

Chesapeake Bay Landcover Data Series (CBLCD). http://landsat.gsfc.nasa.gov/pdf_archive/CBWLCD.pdf

Drummond MA, and TR Loveland. 2010. Land-use pressure and a transition to forest-cover loss in the eastern United States. Bioscience 60: 286-298.

Frenkel A, and M Ashkenazi. 2008. The integrated sprawl index: measuring the urban landscape in Israel. Ann Reg Sci 42:99–121.

Fry J, G Xian, S Jin, J Dewitz, C Homer, L Yang, C Barnes, N Herold, and J Wickham. 2011. Completion of the 2006 National Land Cover Database for the Conterminous United States, *PE&RS*, Vol. 77(9):858-864.

Intergovernmental Panel on Climate Change [IPCC]. 2007. Climate change 2007, Synthesis Report. 107 p. http://www.ipcc.ch/publications_and_data/publications_ipcc_fourth_assessment_report_synthesis_report.htm. [Date accessed: March 25, 2011].

Li X, and P Gong. 2016. Urban growth models: progress and perspective. Science Bulletin 1-16.

Loveland TR, TL Sohl, K Sayler, A Gallant, J Dwyer, JE Vogelmann, and GJ Zylstra. 1999. Land cover Trends: Rates, Causes, and Consequences of Late-Twentieth Century U.S. Land Cover Change. EPA Journal: 1-52.

McGarigal K, Compton BW, Plunkett EB, DeLuca WV, and Grand J. 2017. Designing sustainable landscapes products, including technical documentation and data products. https://scholarworks.umass.edu/designing_sustainable_landscapes/

Nakicenovic N, J Alcamo, G Davis, and others. 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. http://www.grida.no/climate/ipcc/emission/index.htm. [Date accessed: March 25, 2011].

Sleeter BM, and TL Sohl. 2010. The Future of Land-Use in the United States: Downscaling SRES Emission Scenarios. American Geophysical Union, Fall Meeting 2010, abstract #H52A-05. http://adsabs.harvard.edu/abs/2010AGUFM.H52A..05S

Sohl TL, KL Sayler, MA Drummon, and TR Loveland. 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, and K Sayler. 2008. Using the FORE-SCE model to project land-cover change in the southeastern United States. USGS Staff -- Published Research. Paper 572. http://digitalcommons.unl.edu/usgsstaffpub/572

Theobald D. 2005. Landscape patterns of exurban growth in the USA from 1980 to 2020. Ecology and Society 10(1): 32. [online] URL: http://www.ecologyandsociety.org/vol10/iss1/art32/

U.S. Department of Agriculture. 2009. Summary Report: 2007 National Resources Inventory, Natural Resources Conservation Service, Washington, DC, and Center for Survey Statistics and Methodology, Iowa State University, Ames, Iowa. 123 pages. http://www.nrcs.usda.gov/technical/NRI/2007/2007_NRI_Summary.pdf

U.S. Environmental Protection Agency (EPA). 2009. Land-Use Scenarios: National-Scale Housing-Density Scenarios Consistent with Climate Change Storylines. Global Change Research Program, National Center for Environmental Assessment, Washington, DC; EPA/600/R-08/076F. Available from: National Technical Information Service, Springfield, VA, and online at http://www.epa.gov/ncea.

Verburg PH, W Soepboer, R Limpiada, MVO Espaldon, M Sharifa, and A Veldkamp. 2002. Land-use change modelling at the regional scale: the CLUE-S model, Environmental Management, 30: 391–405.

Wear DN. 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.

Wickham JD, SV Stehman, JA Fry, JH Smih, and CG Homer. 2010. Thematic accuracy of the NLCD 2001 land cover for the conterminous United States. Remote Sensing and Environment 114(6):1286-1296.

White R, and G Engelen. 2000. High-resolution integrated modelling of the spatial dynamics of urban and regional systems. Computers, Environment and Urban Systems 24: 383-400.

Zhao Y, and Y Murayama. 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.

10 Appendix A

Table 1. Final logistic regression models and associated statistics for each of the six development transition types for each of the model points in the two-dimensional state space depicted in **Figures 2 and 3**. Note, models were not fit for transitions and model points with <200 training points in the training windows located <2 standard deviations from the model point in the state space (see text for details).

Transition	Average D^2	Average Kappa	Average cross validated Kappa	Average sample size	Number of training windows with models
1 New Low Intensity	0.51	0.75	0.65	233	799
2 New Medium Intensity	0.55	0.78	0.69	206	480
3 New High Intensity	0.59	0.81	0.72	164	304
4 Low to Medium	0.31	0.61	0.51	106	212
5 Low to High	0.42	0.71	0.61	99	180
6 Medium to High	0.33	0.63	0.51	92	122