Practical Activity-Based Model Development and Deployment: Technical Aspects and Field Validation

Ramachandran Balakrishna*

Principal Transportation Scientist

Caliper Corporation

rama@caliper.com

Srinivasan Sundaram

Principal Transportation Engineer

Caliper Corporation

srini@caliper.com

Howard Slavin

President

Caliper Corporation

howard@caliper.com

*Corresponding author

Total page count: 20

Submission date: August 1, 2026

Structured abstract

Objectives

The rollout of Activity-Based Models (ABM) has been slowed by challenges of data availability, model complexity, running time, exacting hardware requirements, and a lack of proper validation. Our objective is to develop a behaviorally sound ABM framework derived from observed patterns in household survey data, that is internally consistent, captures the majority of the decision patterns observed in the survey, and runs fast enough to be a meaningful tool for planning agencies.

Methods

We build our ABM by customizing the framework to local data through model estimation rather than model transfer. Run times are minimized through efficient multi-threading and software engineering.

Findings

A parsimonious ABM with adequate behavioral mechanisms is possible. Such an ABM, when implemented efficiently, can run in hours rather than days and on readily available hardware, allowing agencies to run all necessary scenarios and conduct tests to gain a deep understanding of their model's sensitivities.

Novelty

Our innovation is in building complex ABMs from the ground up against local data, rather than transferring an entire donor model from a different region. We further demonstrate that targeted model transfers can be a useful mechanism though key components must be re-estimated with a local survey. We present a proof-of-concept for the novel use of Big Data for ABM calibration. Our methodology facilitates deploying simpler ABMs when added complexity is not warranted, which is expected to make ABMs more attractive to small and medium agencies with simpler transportation behaviors and limited resources.

Practical Applications

We present three different deployments of the proposed ABM framework to demonstrate its ability to handle a variety of regional challenges, validate its general and flexible architecture, and empirically verify its practical run times despite applying a fully disaggregate treatment to resident demand forecasting.

Introduction

Travel demand models allow public agencies to quantitatively evaluate project benefits and recommend a selection of investments that are likely to best serve the region's future mobility needs. Fully-disaggregate Activity-Based Models (ABM) have recently become the gold standard for high-fidelity demand forecasting. At a high level, most ABMs adopt a largely similar framework based on a sequence of individual decisions typically progressing from long-range to medium-range to short-range. Long-range choices include lifestyle decisions like home location, access to (or dependency on) public transit, household auto ownership, driver license ownership, habitual primary tour mode, etc. Medium-range decisions include work flexibility, full-time vs. part-time work status, child-related escort activity participation, etc. Short-term decisions pertain to day-to-day selections such as mode switches, intermediate stops and non-mandatory activity participation (both individually as well as in groups of household members). The outcomes of this series of choices are then resolved into a daily schedule that must respect mandatory (e.g. work, school) time constraints and available free time budgets for discretionary activities.

ABMs typically apply an analytical model for each of the above decisions, while some platforms choose to directly apply real-world distributions for select model steps (for example, the frequency of intermediate stops by purpose could reflect the frequency of such stops in the survey data). The selection of either approach determines calibration burden and transferability to future scenarios. An analytical model well-calibrated against observed data is expected to predict well with future changes in input variables. If a model is instead based on an observed distribution, it is assumed that a similar distribution will be available for future scenarios as external inputs to the model (or that the current distribution will continue to apply in the future). Given the sparse state of the data typically available for model development, either approach may be deemed appropriate, with the above-mentioned caveats.

The adoption of ABMs in practice has faced challenges relating to data availability and quality, growing model complexity, significant hardware requirements, long run times (in the order of days), and a lack of focus on validating the models against external data. In this paper, we present CCABM, a parsimonious yet behaviorally driven ABM framework and software implementation that resolves all of the above limitations, paving the way for regions of all sizes to deploy ABMs and benefit from their potential advantages.

The paper is structured as follows. We present a review of many of the prominent ABM platforms in use today, through a summary of the most pressing issues challenging the widespread deployment of ABMs today. We then provide a detailed overview of CCABM, with in-depth discussions of its key innovations and modeling advantages. We subsequently describe CCABM deployments in three behaviorally varied regions, together with calibration and validation metrics, and conclude with notes on ABM output stochasticity, lessons learned and future directions.

Literature review

Several ABM methodologies have been developed and deployed to date. These include DAYSIM (Bowman, 2014), SimMobility (Adnan et al., 2016), TASHA (Roorda and Miller, 2006), CEMDAP (Bhat et al., 2004), CT-RAMP (Davidson et al., 2010) and ActivitySim (Zephyr Foundation, 2026). Each platform makes a series of assumptions that, taken together, attempt to replicate observed travel patterns and capture the behavioral processes driving them. A review of the published literature, including project reports, reveals the following essential characteristics of current ABMs:

- Model complexity and resulting lack of sufficient validation
  The model specifications in most current ABMs are significantly complex, and tailor to the policy evaluation needs of a small number of locations. For example, an elaborate mode choice model with numerous public transit modes arranged in several levels of nesting is unlikely to be applicable to large swathes of the United States where public transit service is very limited, spatially restricted, or non-existent. There is no evidence that greater explanatory or predictive power always results from this complexity, especially when it prevents meaningful validations, comparisons to prior, simpler models, or evaluations of model stochasticity owing to excessive run times.

- Expensive minimum hardware requirements
  Some of the most advanced ABMs are known to require computers with up to 1 TB of physical memory and generate upward of 500 GB of intermediate files. The primary cause for this escalated need is often a highly detailed representation of public transit, with resulting skim matrices that escalate running time and memory needs. While less resources are likely to be required for smaller regions, there is space for leaner, parsimonious models that can capture most of the observed behaviors and patterns at a fraction of the compute.

- Internally inconsistent calculations
  ABMs attempt to replicate the complexities of daily schedule planning across millions of residents. The sequential nature of these frameworks often results in misalignment between upstream and downstream model outcomes. One example seen in practice is the need for children's school mode choice decisions to be adjusted post-facto because of the unavailability of an adult to drive them to school. Similarly, independently simulated mode choices across workers in an auto-insufficient household will necessitate post-model corrections. Such actions complicate model calibration by overriding the results of a calibrated step.

- Impractical running times
  Despite the use of very powerful and resource-heavy hardware, agencies continue to report very high run times in practice, often twice or three times the run times of their much simpler trip-based models. This could serve as a barrier to building trust in ABMs through routine use.

These practical realities indicate a pressing need for a tractable and useful ABM methodology and implementation that can accelerate the deployment and regular use of ABMs. We present the design for an ABM that is adequately powerful for small and medium regions, yet has the potential to scale to larger regions, always maintains consistency across all simulated decisions, efficiently runs on readily available hardware, and facilitates practical model validation.

ABM methodology

ABMs generally operate with a few key sequential components encased in a feedback loop (Figure 1):

- Synthetic population generator

- Skim and accessibility calculator
- ABM for resident population
- Supplemental model steps for non-residents
- Demand post-processor
- Network assignment models

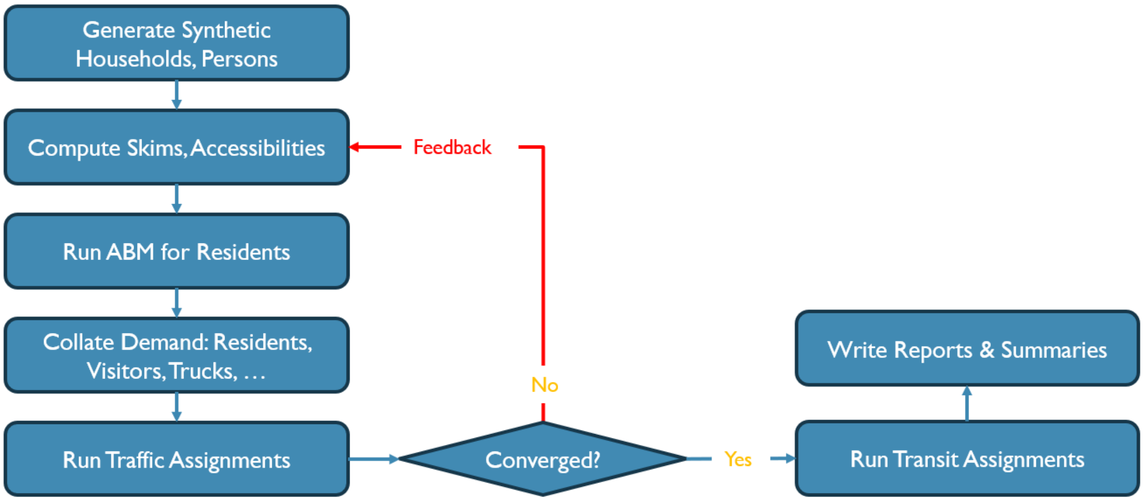

Figure 1: Conceptual Travel Demand Model

The synthetic population generator simulates a full set of the region's residents, together with key demographic and household variables. The skim and accessibility calculator computes important spatial variables such as origin-to-destination travel times and the number of jobs and retail service locations accessible from each origin by the different modes of travel. The ABM step first forecasts the core structure of each resident household's activities, indicating whether each member participates in mandatory, non-mandatory solo, and joint activities. This core is then flesh out with finer details such as activity frequency, time of day, duration, destination and mode. Supplemental model steps integrate non-resident aspects of the region's demand, such as external trips that enter or exit the region from outside its boundary; visitors; and trucks. The demand post-processor collates both resident and supplemental demand by mode and by time of day to assemble the demand inputs for the network assignment models that determine level of service. This sequence of steps is run in feedback to solve the fixed-point problem resulting from not knowing the congested travel times *a priori*. Various ABM frameworks are possible. The next sub-section outlines our framework, named CCABM.

Model framework

The CCABM modeling framework is illustrated in Figure 2 and leverages the Logit family of discrete choice models at every step. The probabilities generated by each Logit application are realized through Monte Carlo simulations, so that an individual or household choice is recorded before the next decision is taken up.

A household's decisions are classified into long-term and short-term choices, with the long-term decisions relating to driver license holding, vehicle ownership and work flexibility. These decisions tend to be stable over time and are informed primarily by the demographic variables predicted for the synthetic population. Work flexibility is particularly relevant due to the current levels of remote and hybrid work schedules maintained by a significant fraction of workers, in addition to part-time work status. Each worker's work industry is forecasted by the population synthesizer, and this allows CCABM to appropriately apply relevant work-from-home and hybrid probabilities individually.

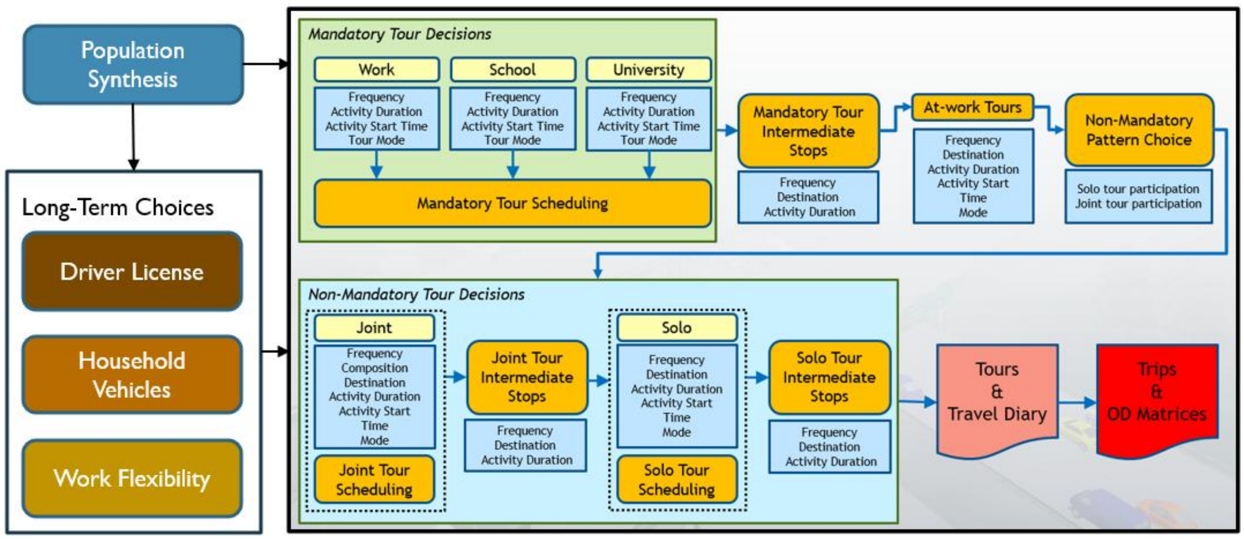

Figure 2: Detailed CCABM framework

Once the long-term decisions are predicted, CCABM moves sequentially through various short-term (or within-day) choices so that mandatory activities (work, school, college/university) are prioritized first. Each household's **daily activity pattern** evolves as details are added to the mandatory pattern, first through intermediate stops and then with the accommodation of discretionary (i.e. non-mandatory) activities. Note that the latter category of activities can be solo or joint, with multiple family members participating in joint activities. As the framework indicates, prior decisions may impact downstream choice alternatives, and care must be taken to ensure that the choice sets remain feasible at all times. The realism of the simulation is further enhanced by adding at-work sub-tours that capture activities that are executed from within the overall work window (e.g. lunch or errands).

The model structure described above is based on basic assumptions about the importance of recurrent mandatory activities over more flexible discretionary ones. However, it should be noted that the framework may be easily revised to reorder certain steps or introduce new steps to capture future behaviors.

The primary output of the ABM calculations are highly-detailed tour and trip diaries that mimic a household travel survey, but span the entire synthetic population. This disaggregate output may be sliced or aggregated on demand to facilitate a range of analyses and policy evaluations. The simplest aggregation remains the translation of the trip diary into time-dependent origin-destination (OD) matrices for input to the relevant network assignment routines. It should be pointed out that the trips are associated with departure times and activity durations to the minute (based on observed time profiles in the survey trip database), so that sequences of trips may also be directly simulated in appropriate Dynamic Traffic Assignment (DTA) platforms.

Large-scale simulations of the interconnected kind described above involve several nuances that are captured in CCABM to improve realism and accuracy. These critical aspects are outlined in the following sub-sections.

Population synthesis

Advanced travel demand models use a synthetic population as the foundation upon which household- and person-level decisions are modeled. A synthetic population is an enumeration of every single resident in the region, together with key demographic information about their household and personal circumstances. It allows the model to predict high-fidelity outcomes that account for disaggregate variables such as age, gender, and work industry, as well as tune the model's predictions to intra-household constraints relating to carpooling, vehicle allocation and joint activity participation. A population synthesizer operates on an area geographic layer (such as block groups or Traffic Analysis Zones (TAZ)), predicts the number of households of each type residing in an area, and conducts weight-based sampling with replacement from matching household records in an input seed household table. Sampled households are transferred into the synthetic population list along with all members in these households. The above synthesis procedure involves two key aspects:

- Determining the number of each household type for every area
- Choosing the sampling weights

Population synthesis is a well-researched topic with a rich body of literature. The reader is referred to Rich et al. (2019) and Ramadan and Sisiopiku (2019) for in-depth coverage of how existing methods handle the above two tasks with differing implications on accuracy and compute burden. Our own review identified the following key requirements that have been implemented for CCABM:

- A method that matches control totals at both the household and person levels
- An algorithm that is robust with the quality and coverage of current travel surveys
- A solution algorithm that runs fast on readily available computing hardware
- An incremental synthesis feature that updates only part of the region's population

Population synthesis in CCABM is based on Iterative Proportional Updating (IPU) (Ye et al. (2009), Konduri et al.(2016)) that extends the popular Iterative Proportional Fitting (IPF) method to include person marginals. IPU maintains the prevalent approach of using IPF to determine the number of each household type needed in every area, but adjusts the seed household weights for each area to allow for a better match to the person marginals. Moving from fixed weights to area-specific weights improves the spatial representativeness of the population along every marginal dimension, which automatically enhances the accuracy of the travel demand model as a whole. CCABM employs a simplification to Konduri et al. by focusing on one-way marginals instead of the jointly distributed marginals. See Balakrishna et al. (2019) and Balakrishna and Sundaram (2019) for a comprehensive analysis of the two IPU variants, including the data and numerical challenges posed by the original method proposed by Konduri et al. and the computational speed-ups gained by our modification. Population synthesis without and with IPU are illustrated conceptually in Figure 3.

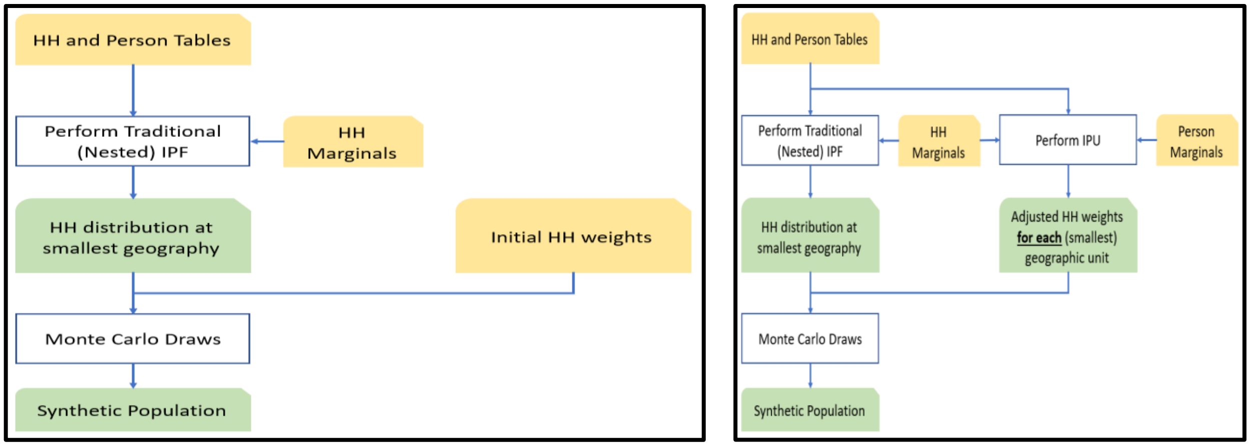

Figure 3: Population synthesis (a) without IPU, and (b) with IPU

Through a sampling of the many marginals that feature in a population synthesis run, Figures 4 and 5 highlight the accuracy added by employing IPU. These charts are particularly revealing, since the youngest and oldest age groups are often associated with significant escort-related travel.

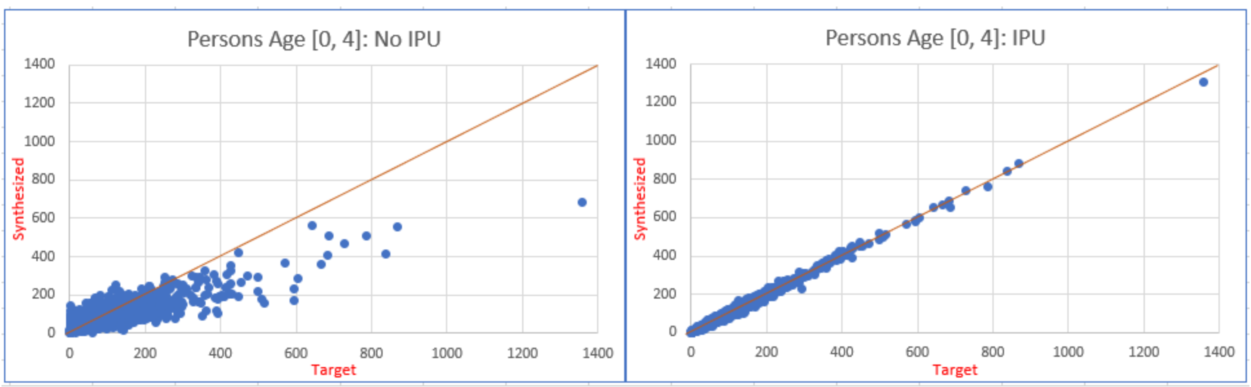

Figure 4: Block group population for age 0-4 years (a) without IPU (b) with IPU

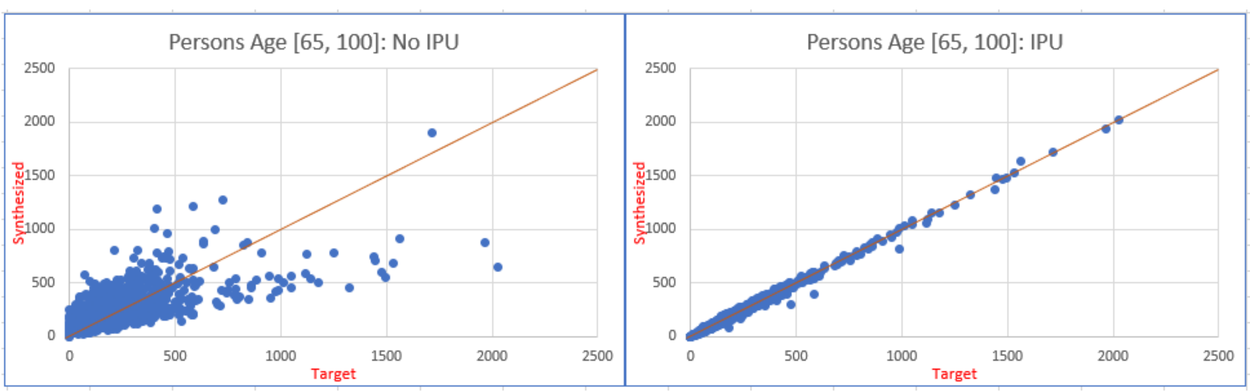

Figure 5: Block group population for age 65-100 years (a) without IPU (b) with IPU

In addition to the above resident synthesis, CCABM generates 1-person "households" to represent college and university dorm residents. These records are based on group quarters numbers for the zones with dorms. Apart from fixing their home locations to the appropriate college/university's zone, individual students' income categories and access to a vehicle are simulated using locally sourced proportions for the general student population in the region. This allows the model to capture non-class-related activities happening in and around campus. It should be noted that college students residing independently are treated separately as part of resident synthesis, and make decisions such as college location choice, mode and time of day decisions for college activities.

Auto ownership decisions

Travel demand models have two broad options for handling auto ownership. The first is to forecast this critical household attribute directly with the synthetic population. This approach will provide the best fit to the input control totals, though auto ownership for all scenarios will depend on the forecasting accuracy of future auto ownership control totals. Further, fixing the household vehicles in the population does not allow the model to capture long-term vehicle holding sensitivities to infrastructural improvements such as added or enhanced public transit service or bike and walk paths.

The second option, adopted by CCABM, is to develop an auto ownership model that predicts each household's vehicle fleet size. This model is sensitive to various demographic effects such as the number of licensed adults, workers, seniors and children; and household income. The decision also accounts for accessibility effects through public transit and bike/walk destination choice logsums interacted with income, and a range of 4D variables as indicated in Table 1.

Beta	**Coefficient**	**t statistic**	**Household Auto Ownership Alternatives**				
			None	**One**	**Two**	**Three**	**Four+**
ASC_None	-3.9929	-7.10	X				
ASC_One	-0.0305	-0.34		X			
ASC_Two	1.4639	17.37			X		
ASC_Three	0.6886	11.33				X	
One licensed adult	2.5696	30.61		X			
Two licensed adults	0.3127	2.88			X		
Three licensed adults	1.1493	11.16				X	
Four+ licensed adults	2.2761	13.13					X
Two adults	0.4207	4.25			X		
HH has only 1 worker	0.3344	4.50		X			
No kids in HH	0.6075	1.51	X				
Kids per licensed adult	0.2229	2.44		X	X		
HH Size >= 5	0.5040	3.08				X	X
HH income category >= 7	0.5124	7.92				X	X
Low-income HH with only seniors	0.4344	1.48	X				
Connectivity Index > 0.5	0.4182	1.14	X				
HH density < 1000	-0.5902	-1.49	X				
Jobs-Housing diversity > 0.5	0.2448	1.09	X				
Transit accessibility to jobs > 1000	0.9947	2.45	X				
Transit accessibility to retail > 1000	0.1230	0.51	X				
Walk index > 0.8	0.2188	0.92	X				
(Income category <= 4) * Non-motorized accessibility	0.1057	3.89	X				
(Income category <= 4) * Public transit accessibility	0.0655	11.82	X	X			
Rho^2	**0.377**						

Table 1: Household auto ownership model specification

Households' decisions on auto ownership thus have the opportunity to evolve with the network and its performance, as feedback iterations update public transit and non-motorized accessibilities, empowering the model with a critical policy lever for congestion management and sustainability.

Nested destination choice

Location decisions drive traffic outcomes. While the home locations are forecasted as part of population synthesis, a variety of other locations must be determined within the ABM. Disaggregate models have the option of using Logit-based destination choice models to capture the behavior underlying location decisions. The estimation of Multinomial Logit (MNL)-based destination choice models, while prevalent in practice, is limited by sparse survey coverage: only a handful of the thousands of zones in the region is likely to be reported as destinations in any household travel survey today. This results in models that fit poorly, have very little explanatory power, and often resort to sampling of alternatives. Many practice models also include multiple mathematical transforms of distance in an attempt to boost model fit, but this approach can have the counter-intuitive effect of farther destinations becoming more attractive than those nearer to home (Figure 6). Distance is also congestion-invariant and removes a critical feedback mechanism that can have significant impacts on people's location choices.

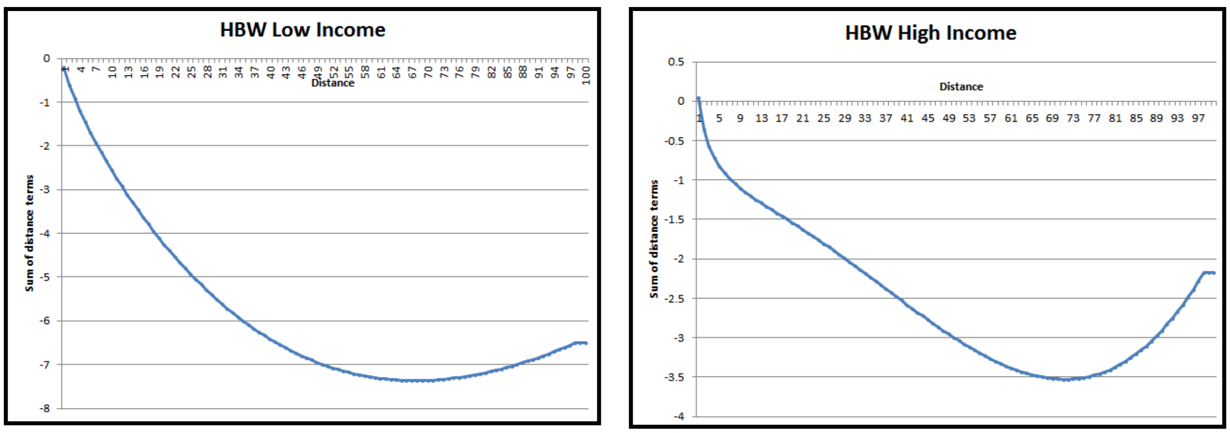

Figure 6: Non-monotonic, counter-intuitive distance effects due to multiple transforms

Destination choice models in CCABM adopt a nested approach that alleviates all the above challenges. Location decisions are hypothesized to happen in two stages. In the first stage, people choose a general part of the region they would like to travel to. This is facilitated by partitioning the region's zones into non-overlapping clusters based on existing features such as towns, districts and counties. In the second stage, a zone is picked conditional on a chosen cluster. An example of the corresponding is illustrated in Figure 7.

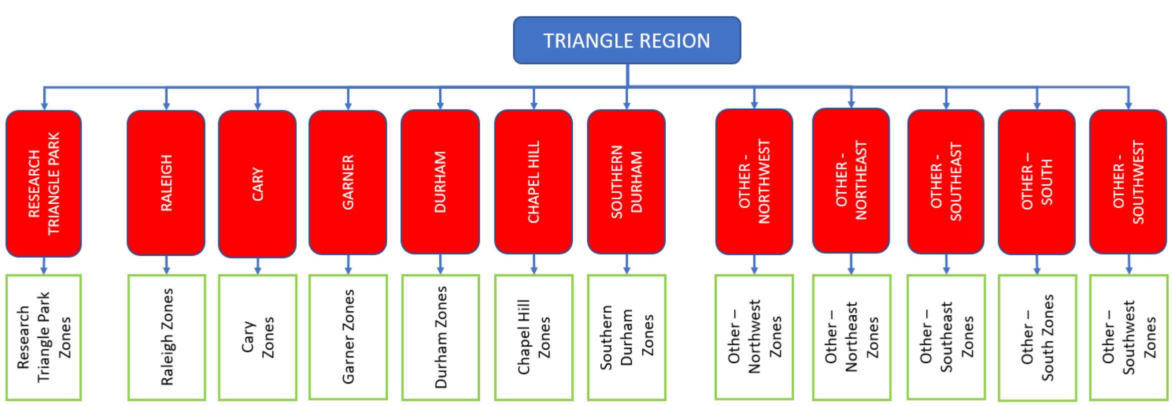

Figure 7: Nested Destination Choice in the Triangle model (North Carolina)

The nested approach has many advantages. First, surveys will almost always provide full and meaningful coverage at the cluster level, thus improving the reliability of model estimation. Trip patterns and trip lengths are more accurately predicted since errors are now localized to the correct clusters. Second, it allows for specifications that capture the overall value derived from choosing each cluster, including variables that capture the range of attractions and opportunities available at the cluster level. Third, the behaviorally grounded approach allows for testing more useful variables such as travel times and mode choice logsums in place of fixed distance effects. Finally, the model estimation will automatically reduce the specification to MNL when needed by estimating nest coefficients that are not statistically different from 1. While nested destination choice models may be estimated as a sequence of MNL models, the more efficient simultaneous estimation is also feasible with open-source packages like Larch (Newman, 2026) and is the method adopted for CCABM.

A sample work destination choice specification estimated for Monterey Bay, California is presented in Tables 2 and 3 to highlight the person-level fidelity and model fit achieved by CCABM.

Variable	Coefficient	t-Statistic
Constant (NorthWest)	0.5189	1.44
Constant (North Vinyards)		
Constant (Watsonville)	0.8898	2.55
Constant (Gabilan Acres)		
Constant (Salinas)	1.1733	3.95
Constant (Santa Cruz)	1.9776	8.58
Constant (Hollister)	0.7739	2.42
Constant (Spreckels Chualar Gonzales)	0.4541	2.94
Constant (Del Monte Marina)	1.0903	3.56
Constant (SouthEast)	1.2399	2.22
Constant (North Salinas)	0.9218	3.11
Constant (SouthWest)	1.0288	4.96
Constant (Capitola Rio Del Mar)	1.3755	5.28
Constant (Monterey Carmel)	1.7746	6.13
Constant (Moss Landing)		
Constant (Soledad Greenfield King)		
Non-Motorized MC Logsum (v0 or vi)	0.1065	3.06
Intra-zonal	1.5617	12.33
Intra-cluster	0.2530	4.26
Drive time	-0.0760	-25.18
Drive time in excess of 60 minutes	0.0270	2.56
Distance > 10 miles	-0.1368	-1.79
Size coefficient	0.6958	-9.45
Rho^2	**0.162**	
Size term is Size * ln(1 + Employment)		

Table 2: Utility specification for work location choice for Monterey, California

Cluster	Nest Coefficient	t-Statistic
NorthWest	1.000	
North Vinyards	1.000	
Watsonville	0.829	-2.22
Gabilan Acres	1.000	
Salinas	0.821	-2.77
Santa Cruz	0.701	-7.00
Hollister	0.860	-2.21
Spreckels Chualar Gonzales	1.000	
Del Monte Marina	0.942	-0.94
SouthEast	1.000	
North Salinas	0.781	-3.26
SouthWest	1.000	
Capitola Rio Del Mar	0.722	-5.56
Monterey Carmel	0.777	-3.65
Moss Landing	0.945	-1.16
Soledad Greenfield King	1.000	

Table 3: Nest coefficients for work location choice for Monterey, California

Intra-household interactions

While households in the real world must reconcile their own unique constraints and desires daily, CCABM pays special attention to the most common of these intra-household interactions that generalize across a majority of the households. The first such interaction involves the **escorting of children** to and from school. Dubbed PUDO (Pick-Up/Drop-Off) in the transportation literature, this component critically influences adult household members' daily schedules. CCABM handles PUDO through two mechanisms that ensure consistency across all decisions made by both children and adults:

- A child's school mode choice decision will include the PUDO option only if there is a qualified adult to cover this activity. A hierarchical priority order checks for non-working adult(s) with a driver license, followed by working adult(s) with a work window flexible enough to accommodate PUDO.
- If PUDO is chosen with a qualified working adult, the adult can either execute a separate PUDO tour or cover PUDO as a stop on their way to/from work (which might further trigger a work schedule adjustment).

If PUDO is not deemed to be an option, then the child must choose from other modes including walk, bike and riding the school bus (driving alone to school could be an option for the older high school children). The above logic ensures consistency between adult and child decisions and eliminates the need for post-model adjustments to previously simulated outcomes.

The second interaction involves the **utilization of vehicles** available to each household. This is particularly relevant in auto-insufficient households where the number of working adults with a driver license exceeds the number of vehicles held by the household. In keeping with the behavioral assumption that a household's fleet of vehicles is purchased with regular long-term use in mind, CCABM will maximize fleet utilization within each household through a tiered vehicle allocation process that prioritizes, full-time workers, part-time workers, university students and non-working adults in that order.

After each adult makes their mode choice decision for the day's mandatory activity, the available vehicles are scanned to ensure that unused vehicles are immediately made available to the next tier.

A third dimension to intra-household interactions is the compensatory relationship between **intermediate stops and separate non-mandatory tours**. CCABM recognizes that people can satisfy their non-mandatory activities either by making stops along mandatory tours, or by executing separate tours that include no mandatory activities. Making more of the first will reduce the need for the second, and vice versa. This trade-off must be calibrated against the local travel survey to ensure appropriate splits, and gives CCABM the power to predict, for example, that significantly higher congestion in a future scenario might cause residents to make more intermediate stops and thus fewer separate non-mandatory tours, thereby reducing overall vehicle miles travelled (VMT).

Finally, CCABM uses a **time manager** object to track each individual's time use as the model moves through its steps. This persistent object plays a crucial role in ensuring model consistency throughout the process, as it might be queried at any time to ascertain currently available free time for all individuals. This knowledge is essential for considering additional non-mandatory activities in the day. The time manager is further able to report on overlapping free time across multiple members of a household, which is relevant for joint activity scheduling. A separate **scheduler object** is tasked with ensuring that mode-specific travel times are considered while estimating tour departure time periods.

Tour mode choice

It is important that the ABM realistically capture residents' trade-offs when choosing travel mode, as these decisions collectively determine network loadings and congestion levels. The mode choice component is particularly relevant in regions with meaningfully competing mode alternatives, and especially with public transit. The level of complexity in the mode choice model specifications must depend on the complexity of options on the ground, so that transit-rich regions might benefit from Nested Logit (NL) structures that capture both compensatory and competing relationships between drive, public transit and non-motorized modes. Regions that are mostly auto-focused should use simpler specifications (including MNL structures) and include only the options that are likely to be part of policy analyses in the near term. CCABM can flexibly handle both MNL and NL mode choice components, and use a mix of the two across activity purposes. The work tour mode choice model for Oahu, Hawaii is shown in Table 4.

Beta	Unit	Coefficient	t statistic	Mode Alternatives						
				SOV	HOV2	HOV3+	Bus	Walk	Bike	TNC
Constant (SOV)		3.0500	2.5	X						
Constant (HOV2)		1.6000	1.2		X					
Constant (HOV3+)		0.7990	0.7			X				
Constant (Bus)		3.1900	2.4				X			
Constant (Walk)		5.5500	4.8					X		
Constant (Bike)		2.6700	2.2						X	
Drive time (Auto modes)	Minutes	-0.0912	-3.3	X	X	X				
Drive time (TNC)	Minutes	-0.2080	-2.2							X
Total public transit time	Minutes	-0.0267	-2.7				X			
Walk distance	Miles	-1.9400	-8.6					X		
Bike distance	Miles	-0.4370	-2.7						X	
Public transit fare	$	-0.6060	-0.9				X			
Access + Egress time less than 20 minutes		0.3770	1.1				X			
Origin and Destination transit stop density > 50		0.3280	1.0				X			
HH has vehicle(s) (SOV)		2.7300	6.4	X						
HH has vehicle(s) (Carpool)		2.8000	7.4		X	X				
Female (Carpool)		0.2840	2.6		X	X				
Female (Bus)		0.3040	1.2				X			
Female (TNC)		1.5200	1.4							X
HH income less than $50K		0.2230	0.9				X	X	X	
Rho-squared		**0.424**								

Table 4: Mode choice specification for work tours in Oahu, Hawaii

Time-of-day decisions

Activity timing and duration play a central role in determining the temporal loading of the transportation networks. Apart from helping decide tour departure time (and hence departure time period) conditional on tour mode, these decisions also identify the gaps of free time available to individuals for additional activities. These gaps essentially form constraints for subsequent activity participation models. Behaviorally, activity start times and durations are not independently chosen. CCABM adopts a joint specification for all time-of-day computations, with each Logit alternative comprised of both an activity start period and end period. The activity duration is therefore implicit in this choice, and exact start and end times may be further simulated based on data-driven temporal profiles if needed.

Table 5 illustrates a sample time-of-day specification, for work activities. Figure 8 visualizes the joint distribution of activity start and duration by hour of the day and by time period, showing clear trends in work schedule arrangements.

Beta	Coefficient	t statistic	Forward-Return Time Period Alternatives																				
			EA-EA	EA-AM	EA-MD	EA-PM	EA-EV	EA-NT	AM-AM	AM-MD	AM-PM	AM-EV	AM-NT	MD-MD	MD-PM	MD-EV	MD-NT	PM-PM	PM-EV	PM-NT	EV-EV	EV-NT	NT-NT
ASCs			-1.0724	-0.4074	3.2289	1.8063	0.6243	-1.1249	-0.1220	3.6114	3.2347	1.7501	0.1133	3.1837	3.4779	3.1152	1.0599	-0.3048	0.6952	0.9740	0.1426	-0.1451	
ASC t-statistics			-1.31	-0.63	7.52	3.85	1.25	-1.67	-0.19	8.17	7.03	3.71	-0.21	7.15	7.85	7.26	2.22	-0.49	1.38	2.03	0.26	-0.25	
Auto time < 30 min	0.1749	1.17								X				X	X			X	X		X	X	
Auto time between 30 and 60 min	0.1712	1.11								X	X				X	X			X	X	X	X	
Auto time >= 60 min	0.9011	1.94		X	X																		
Full-time worker	1.4462	7.50				X	X	X			X	X	X				X						
Female with kid(s)	0.3044	1.90								X	X			X	X								
High income HH	0.2779	2.31									X	X											
Low income HH	1.1881	4.15	X	X														X	X	X	X	X	
Makes university tour(s)	1.0707	0.94			X					X					X	X	X		X	X			
Makes 2+ work tours	3.0359	8.75	X	X	X				X	X				X	X	X	X	X	X	X	X	X	
Sample size	1428																						
Rho^2	0.278																						

Table 5: Work activity start and end time model specification

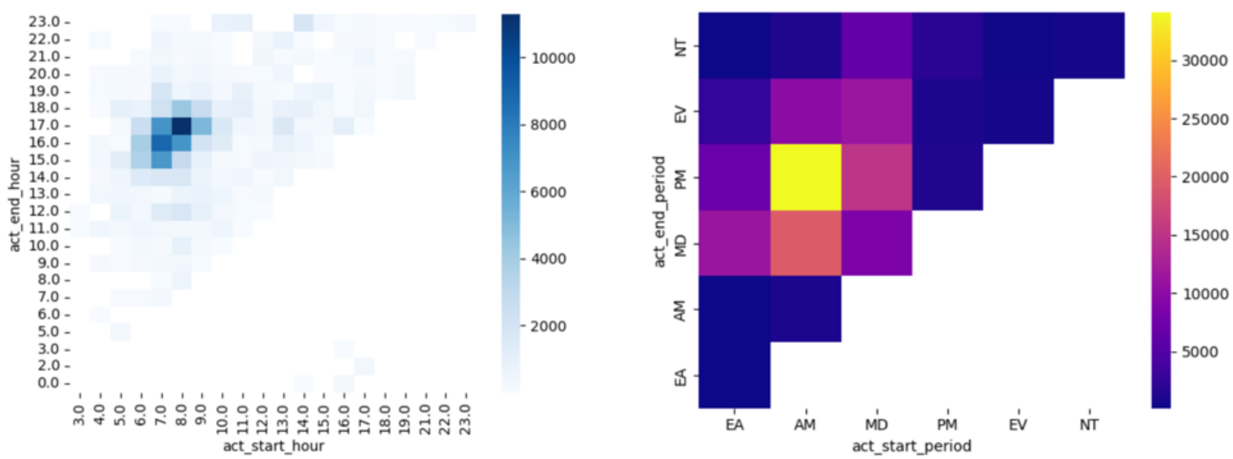

Figure 8: Work time of day distribution by (a) hour (b) time period

Joint tour participation

Joint tours have one additional decision that is absent from all other choices in the ABM: the subset of household members that participate in each joint activity. Each travel party is made up of some combination of adults and children. The travel survey reveals the dominant combinations and maximum party size in the region, covering a variety of activity purposes including shopping, recreation, personal business, errands and dining. The alternative corresponding to the entire household participating together may be included in the choice set. CCABM uses a Logit model that predicts the general combination of adults and children for each joint activity. This is followed by individual-specific models that evaluate the probability of each person participating in a joint activity. The two models in tandem, run sequentially, yield specific instances of household member combinations for the joint activities previously predicted.

Assignments, feedback iterations and convergence

The ABM calculations described thus far pertain only to the demand side of the travel demand model. Traffic forecasts and analyses also require complementary supply side models that determine level of service by mode. While this paper primarily focuses on the demand innovations for practical ABMs, a note on the supply models is necessary. CCABM leans on tested, state-of-the-art traffic and transit assignments already deployed in practice with trip-based models. The detailed outputs from CCABM are therefore aggregated into time-period-specific OD matrices for network assignment. The reader is referred to Slavin et al. (2015) for a detailed review of traffic assignment in practice, together with practical guidance on minimum convergence thresholds for reliable project evaluation. Congested skims from converged assignments are returned to the ABM via feedback to mimic the learning process by which trip-makers gather knowledge about network conditions and alternate routes, and iteratively work toward the best perceived paths.

As noted previously, other assignment methods are also possible. The trip diary from CCABM may be directly simulated with Dynamic Traffic Assignment to exploit the fine temporal resolution of the trip start times, as well as other high-fidelity outputs including continuous value-of-time (VOT) that are more appropriate for simulation based assignment than traditional static assignment.

Computational efficiency

A key requirement for CCABM was that the model should run fast enough to facilitate rapid scenario evaluation with adequate feedback loops. CCABM achieves this via a highly efficient software implementation that automatically multi-threads as many computations as possible to fully exploit the available computing resources. One full pass through an iteration of the ABM, together with the necessary assignments, takes between 1 and 2 hours depending on the scale of the network and demand, which translates to full model runs in the order of a relatively small number of hours. Such parsimonious run times are achieved on readily-available hardware with not more than 64 GB of RAM.

Deployment studies

In this section, we provide details about three practical deployments of CCABM and highlight the unique modeling circumstances in each instance. Table 6 outlines the scale of the three networks.

Region	Highway Nodes	Highway Links	Zones	Households	Population
Monterey Bay, California	19,797	28,320	1,673	253,946	733,673
Peoria, Illinois	4,434	8,790	668	186,058	457,743
Oahu, Hawaii	23,762	30,483	1,117	337,106	983,411

Table 6: Descriptions of deployment regions

Monterey Bay, California

The original deployment of CCABM was for Association of Monterey Bay Area Governments (AMBAG), who supported its development with local data, funding and a partnership with Caltrans and two other agencies: Santa Barbara Council of Area Governments (SBCAG) and San Luis Obispo Council of Governments (SLOCOG). Unique to this effort was the data pooling that allowed for the robust estimation of every model step in CCABM. In addition to combining the California Household Travel Survey datasets from the three agencies, the relevant records from the National Household Travel Survey were also rolled in and re-weighted to double the sample size. The common ABM framework estimated through this process was then calibrated for AMBAG to create a stand-alone model for the agency.

Each model component was individually calibrated to match target shares derived from the weighted survey, taking care to process the models in the right sequence so that each component was calibrated based on a fully calibrated set of components before it. The standard Root Mean Square Error (RMSE) metric, computed as a percentage of the average observed traffic count, is reported in Table 7 by link type and by volume group.

$$\%RMSE = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(y_i^{Obs} - y_i^{Mod}\right)^2}}{\frac{1}{N}\sum_{i=1}^{N} y_i^{Obs}} * 100$$

where N is the number of links with traffic counts; y_i^{Obs} is the count observed on link i; y_i^{Mod} is the model flow on link i.

Link Type	RMSE (%)
ALL	32.02
Caltrans Counts	23.94
Freeways	17.32
Major Arterials	25.16

Link Volume Group	RMSE (%)
ALL	32.02
[10000, 12500)	29.32
[12500, 15000)	27.99
[15000, 17500)	21.47
[17500, 20000)	20.62
[20000, 25000)	20.29
[25000, 35000)	18.82
[35000, 55000)	12.89

Table 7: Fit to traffic counts for AMBAG (a) by link type (b) by volume group

Model validation was performed by comparing the calibrated base year model's congested travel time estimates for the peak periods against the range of times reported by Google for a representative selection of OD pairs. Table 8 summarizes the validation comparisons and confirms the ability of the calibrated CCABM to replicate real-world congestion patterns.

Origin	Destination	Time of Day	Time (Minutes)	
			Model	Google
Carmel by the Sea	Santa Cruz	AM	58.5	50-70
		PM	68.5	55-80
North-West	Santa Cruz	AM	33.2	30-45
		PM	33.2	30-45
South-East	North-West	AM	165.9	140-170
		PM	173.4	140-170
Hollister	Santa Cruz	AM	69.0	50-80
Santa Cruz	Hollister	PM	71.3	65-110
Hollister	Carmel by the Sea	AM	57.1	50-70
		PM	55.7	50-70
San Benito	Santa Cruz	AM	102.4	85-110
		PM	103.0	90-120
Del Monte/Marina	Monterey	AM	24.7	26-35
		PM	24.8	28-40

Table 8: AMBAG travel time validation at the OD level

Peoria, Illinois

Peoria presented a unique opportunity to test model transferability, given that the region did not have a recent travel survey. Monterey Bay and Peoria are similar in their medium size and prevalence of agricultural and manufacturing economies. The AMBAG model with its pre-trained parameters was transferred to Peoria and its population synthesis updated to use local seed tables and marginals. Several of the model components were calibrated to shares extracted from Big Data sourced from location-based services (LBS). The model achieved a %RMSE of 29% on the freeways, which might be further improved through an investment in a travel survey for the region. A key takeaway from this deployment exercise is that Big Data can aid model transfer, however it should not be viewed as a substitute for a well-designed local travel survey.

Oahu, Hawaii

A second model transfer experiment was conducted by re-training the AMBAG CCABM for Oahu, another region similar in size to Monterey Bay. Oahu is distinct in several ways, given its island status, unique topography with natural barriers to travel, and an economy with a significant tourism from the mainland US and from abroad. Key CCABM components (destination and mode choice) were re-estimated with the Oahu survey. In addition, its existing visitor demand model component was replaced with a custom-built disaggregate model inspired by the ABM and supported by a visitor survey conducted at the airport. The %RMSE metric is reported in Table 9 by link type and by volume group.

Link Type	RMSE (%)
ALL	36
Arterial	38
Collector	58
Freeway	15
Local	119
MLHighway	26
MajorArterial	37
MajorCollector	72
Ramp	40

Link Volume Group	RMSE (%)
ALL	36
10000	62
25000	43
50000	33
100000	24
100000+	12

Table 9: Fit to traffic counts for Oahu (a) by link type (b) by volume group

A note on model stochasticity

ABMs are often criticized for their stochastic nature: runs made with identical inputs can produce slightly different outcomes. This is a valid concern, as it directly impacts modeling agencies' ability to evaluate project impacts and prioritize infrastructure budgets. Ideally, such decisions must be made based on sufficiently many model runs with the same input data, so that the base year and each scenario are compared as distributions on the metrics of interest. We take the first step toward this paradigm by quantifying the variance of the VMT predicted by CCABM. Table 10 summarizes the mean, standard deviation and coefficient of variation (%) of the VMT for each AMBAG county and for the entire region, computed from twenty runs of the ABM made with the same base-year inputs. The results indicate that the ABM, while stochastic, is highly stable in a key aggregate measure used by all agencies.

Statistic	Monterey	San Benito	Santa Cruz	Full AMBAG Region
Mean	9,738,211	1,647,816	4,668,566	16,054,593
Std, Dev.	8,729	3,866	7,175	10,873
CV (%)	0.09	0.23	0.15	0.07

Table 10: VMT statistics across twenty runs of the AMBAG ABM

Conclusion

We identify significant challenges faced in the travel demand modeling practice on deploying ABMs on a wide scale. The increasing complexity of such models, with resulting run time and hardware impacts that preclude robust validation and a deeper study of model stochasticity, has caused many planning agencies to maintain trip-based models in parallel for official use. Apart from the associated expense, maintaining two separate models limits trust in the ABM's capabilities and prevents full leverage of the strengths of disaggregate travel demand modeling. We bridge this gap for small- and medium-sized agencies by developing an ABM framework that is behaviorally appropriate, estimated on local data, and calibrated to capture observed patterns adequately while simultaneously minimizing its compute footprint. The paper outlines the key considerations that control model accuracy, and presents detailed sample model specifications to underscore the disaggregate nature of the computations. We further demonstrate the practical nature of the proposed framework on three separate and slightly different applications, including the transfer and re-calibration of the model to other geographic regions.

While our work is ground-breaking toward mainstreaming ABM use in practice, there remain challenges to be overcome. ABMs in general will benefit from a careful expansion in the data used to build them. While we demonstrated the use of Big Data in model calibration, it does not yet replace the household travel survey for key behavioral models such as mode and destination choice. Destination choice models in particular continue to rely on a very small set of variables, and research into expanding this set will help improve their ability to explain location choice. CCABM may also be extended to predict trip mode choice and inter-household carpooling, and track the use of each household's vehicles on a fine time horizon. Machine learning (ML) and artificial intelligence (AI) are exciting tools to test for model calibration, CCABM should also be tested on larger regions and more complex mode choice situations.

Acknowledgements

The authors thank Caltrans, Association of Monterey Bay Area Governments, Santa Barbara Council of Area Governments, San Luis Obispo Council of Governments, Tri-County Regional Planning Commission (Peoria), and Oahu MPO for their funding, data, and for feedback.

References

1. Adnan, M., Pereira, F.C., Azevedo, C.M.L., Basak, K., Lovric, M., Feliu, S.R., Zhu, Y., Ferreira, J., Zegras, C. and Ben-Akiva, M.E. (2016) SimMobility: A multi-scale, integrated, agent-based, simulation platform. Presented at the TRB annual meeting.
2. Balakrishna, R. and Sundaram, S. (2019) Practical population synthesis enhancements to support advanced and activity-based travel demand models. World Conference on Transport Research.
3. Balakrishna, R., Sundaram, S. and Lam. J. (2019) An enhanced and efficient population synthesis approach to support advanced travel demand models. Presented at the TRB annual meeting.
4. Bhat, C.R., Guo, J.Y., Srinivasan, S. and Sivakumar, A. (2004) A comprehensive econometric microsimulator for daily activity-travel patterns. Transportation Research Record, Vol. 1894, pp. 57-66.
5. Bowman, J.L. (2014) DaySim. Presented at the Activity-Based Modeling Symposium, Sydney.
6. Davidson, W., Vovsha, P., Freedman, J. and Donnelly, R. (2010) CT-RAMP family of activity-based models. Proceedings of the Australasian Transport Research Forum, Canberra, Australia.
7. Konduri, K.C., You, D., Garikapati, V.M. and Pendyala, R.M. (2016) Enhanced synthetic population generator that accommodates control variables at multiple geographic resolutions. Transportation Research Record, Vol. 2563, Issue 1.
8. Newman, J.P. (2026) Larch: The Logit architecture. URL: https://github.com/driftlesslabs/larch (Accessed: 29 July 2026)
9. Ramadan, O.E. and Sisiopiku, V.P. (2019) A critical review on population synthesis for activity- and agent-based transportation models. In: Transportation Systems Analysis and Assessment. URL: https://www.intechopen.com/chapters/67163 (Accessed 30 July 2026).
10. Rich, J., Flotterod, G., Garrido, S. and Pereira, F. (2019) Review of population synthesis methodologies. Presented at the hEART conference.
11. Roorda, M. and Miller, E.J. (2006) Assessing transportation policy using an activity-based microsimulation model of travel demand. ITE Journal, November.
12. Slavin, H., Lam, J. and Nanduri, K. (2015) Traffic assignment and feedback research to support improved travel forecasting. Project final report prepared for the Federal Transit Administration (FTA).
13. Ye, X., Konduri, K.C., Pendyala, R.M., Sana, B. and Waddell, P. (2009) Methodology to match distributions of both household and person attributes in generation of synthetic populations. TRB annual meeting.
14. Zephyr Foundation (2026) ActivitySim. URL: https://zephyrtransport.org/activities/activity-sim/ (Accessed 30 July 2026).