Learning evolving populations from multi-year data: A self-contained model for short-term synthetic population forecasting

Vianey Darsel*[1], Etienne Côme[1], and Latifa Oukhellou[1]

[1]COSYS-GRETTIA, Gustave Eiffel University, France

Short summary

Whereas synthetic populations constitute a fundamental component of agent-based simulation models, synthetic population forecasting seeks to generate such populations at future time horizons, enabling the simulation of hypothetical scenarios with representative agents. Existing approaches typically rely on static, single-year snapshots augmented with external data sources. In this work, we propose a novel framework for short-term forecasting that captures temporal trends directly from multi-year data, eliminating the need for external inputs. Our approach combines a projected Bayesian Network with a projective Iterative Proportional Fitting procedure. For each component, we perform attribute-level hyperparameter selection to differentiate stable relationships from those exhibiting meaningful temporal dynamics. Experimental results across diverse scenarios including varying training sample sizes, forecast horizons (one to five years), and attribute configurations, demonstrate that the proposed method yields improved distributional accuracy compared to static models, while maintaining comparable levels of realism. Additionally, we introduce a guardrail-based validation mechanism to detect and mitigate erroneous extrapolations.
Keywords: Activity-based modelling, Synthetic population forecasting, Methodology, Population Synthesis.

Our research is fully reproducible, with open-source datasets and code available[1].

1 Introduction & Literature

Agent-based simulation models rely on a synthetic population composed of discrete agents. Using raw microdata directly is often unfeasible due to issues of representativeness and the risk of compromising individual privacy. Consequently, population synthesis aims to generate agents that reflect the characteristics of an original dataset by learning its underlying joint distribution.
Historically, the literature has primarily focused on generating populations for the same spatio-temporal parameters as the source data (Beckman et al., 1996; Sun & Erath, 2015; Borysov et al., 2019). However, the primary utility of transport simulations is to predict the impacts of hypothetical future scenarios or policy modifications, such as the implementation of Low Emission Zones, the construction of new infrastructure, or the introduction of public transit lines. These projects typically span several years, during which the underlying population evolves. As noted by Darsel, Côme, Rich, et al. (2026), a static population becomes increasingly inaccurate over time; its distribution drifts, rendering it less representative of the future scenarios it is intended to simulate.
To bridge this gap, **agent-forecasting** seeks to generate a synthetic population for a future time horizon $(t + n)$. Several techniques currently address this challenge. Typically, these approaches begin by constructing an initial solution that ensures a representative and privacy-preserving population at the time of the observed dataset, using population synthesis methods. This baseline population is then evolved over time through a range of projection strategies:

- **Static Projection**: This approach utilizes a synthesis model that incorporates aggregated statistics (marginals) for the target year (Lomax et al., 2022). Algorithms such as Iterative Proportional Fitting (IPF) are commonly used for this purpose to update an initial joint distribution to fit these future marginals.

[1]`https://github.com/vdarsel/Synthetic-Population-Forecasting_BN-IPF_model`

- **Dynamic Projection**: This method evolves a baseline population year-over-year until the target date is reached (Namazi-Rad et al., 2014; Fatmi & Habib, 2017). This requires simulating demographic life events—such as births, deaths, and marriages—based, through established transition rates.
- **Resampling**: This technique adjusts a population to match specific aggregated statistics (Prédhumeau & Manley, 2023). The population is typically aged, and individuals are subsequently duplicated or removed to align the sample with target totals.
- **Hybrid Methods**: These approaches combine the aforementioned techniques based on data availability. Kukic & Bierlaire (2025) present a year-over-year approach employing Resampling to update the population when data about a specific year is available, and Dynamic Projection otherwise.

Despite their usefulness, these methods exhibit several limitations. Resampling and static projection approaches depend on future aggregated statistics, which are often estimated and therefore prone to error. Dynamic projection methods, on the other hand, require explicitly defined transition rules that may fail to capture complex interdependencies between attributes. Moreover, all of these approaches rely heavily on external data and typically treat attributes in isolation, overlooking the evolution of their joint distribution. Our preliminary study indicates that these methods struggle to maintain accuracy across multiple attributes Darsel, Côme, Rich, et al. (2026), particularly as most existing projection models generally only consider a single year of historical data.

In this paper, we propose a forecasting method to capture multivariate temporal trends for short-term projections (one- to five-year horizon). This approach allows for the sampling of an updated population by leveraging multi-year data without requiring additional external data sources. The proposed model employs a Bayesian Network (BN) projected to the target-year to derive an initial solution, combined with a projective Iterative Proportional Fitting (IPF) step to synthesize the target population. Unlike existing projection techniques, our model leverages by multi-year datasets to autonomously forecast the evolution of the population structure. Both projection models are based on regression, requiring careful hyperparameter calibration and post-verification to avoid erroneous extrapolation. We evaluate our approach using French census data across multiple time horizons and data scenarios, including varying training set sizes and attribute configurations, in order to assess its robustness and generalisation capabilities.

Our contributions are as follows:

- A temporal projection framework extending Bayesian Networks and Iterative Proportional Fitting to multi-year trends.
- A sequential hyperparameter optimization protocol for multi-year datasets.
- A validation approach using realism *guardrails* to detect inconsistent projections.
- A sensitivity analysis of performance across time horizons, attribute complexity, and training size.

2 Methodology

The primary objective is to project the synthetic population for a target year y_P, given historical data from years $y_1, ..., y_m$, which serve as a training set. As established by Rich (2018), population synthesis frameworks typically generate an initial baseline agent that requires refinement through a fitting stage; this process facilitates the incorporation of additional information and improves the overall representativeness of the solution.

Our approach is described in Figure 1. For the first step, to generate the initial solution, a projected Bayesian Network (BN) is used by sampling from evolving joint distributions. This step aims to catch the evolving relationships between attributes. For the second step, a Iterative Proportional Fitting (IPF) refines the synthetic population to incorporate marginal projections in the final synthetic population to ensure that the marginal projections are respected.

Relying only on a projective IPF considering historical training data would result in a weighted replication of the existing samples, which is undesirable from the perspectives of population diversity and privacy preservation.

To calibrate the distribution projections, hyperparameters are required and are variable-specific to learn the individual behaviour of each attribute. Therefore, the model is calibrated using a walk-forward sliding window for validation before testing, as depicted in Figure 3.

In the following subsections, we describe each stage of the model in detail.

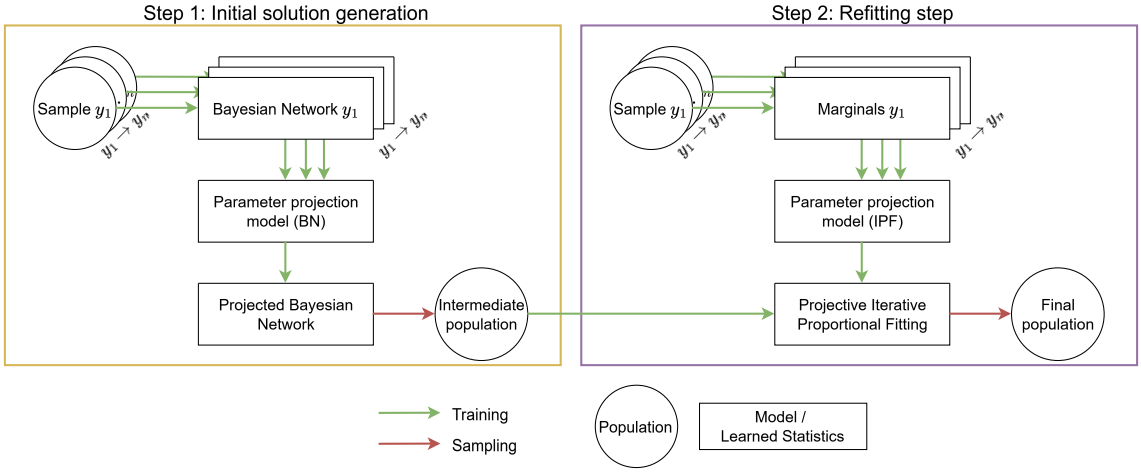


Figure 1: Our projection model is composed of two projected models: one to compute the initial solution, and one for the refitting step.

Projected Bayesian Network

Bayesian Networks (BNs) are Probabilistic Graphical Models (PGMs) defined by a Directed Acyclic Graph (DAG) that encodes the conditional dependencies between variables (Heckerman, 1998). This graph structure facilitates the generation of synthetic samples by decomposing the joint probability distribution into a series of local conditional distributions. Consequently, due to their parsimonious parametrization (in contrast to DGMs), BNs are highly suitable for model projection, enabling direct temporal extrapolation of model coefficients. Furthermore, individuals can be sampled sequentially by following the topological ordering of the DAG. BNs are particularly effective for population synthesis, as they have been shown to outperform other generative methods in capturing attribute dependencies within moderate dimensional datasets (Darsel et al., 2025).
In our model, we assume that the structural dependencies between attributes remain invariant, while the strength and nature of these relationships may evolve over time. Specifically, we employ a BN framework where the DAG describing the population remains constant across all time steps, while the temporal evolution of the population is captured within the Conditional Probability Tables (CPTs), what is illustrated by Figure 2. Under this assumption, while the set of causal or correlative influences between attributes is fixed, the probability distributions governing those influences are time-dependent. This formulation allows the model to generate a synthetic population for any arbitrary target year.
This approach necessitates two distinct learning phases: first, identifying a robust structure that remains valid throughout the temporal horizon—including the target year—and second, modelling the temporal trajectory of the conditional probabilities within that structure.

Fitting the Directed Acyclic Graph (DAG) In our model, the DAG $\mathcal{G}$ is assumed to be temporally invariant. To ensure that the structural dependencies reflect the most contemporary population dynamics, we perform structure learning using the most recent dataset (i.e., for year y_m). In our implementation, $\mathcal{G}$ is identified via a score-based approach to optimize the Bayesian Information Criterion (BIC) using a hill-climbing heuristic (Koller & Friedman, 2009). This optimization facilitates the discovery of complex, non-tree architectures, enabling the model to capture intricate dependencies between multiple attributes.

Projecting the Conditional Probability Tables (CPTs) The primary challenge of the proposed model lies in the temporal projection of the CPTs. Since the DAG structure is assumed to be time-invariant, we estimate the CPTs parameters independently for each year $y_t \in \{y_1, ..., y_m\}$ within the training set. This temporal series of parameters enables the modelling of the temporal evolution of the conditional probabilities.
Because probability coefficients $c_{i,t}$ (coefficient i for year y_t) are constrained to the interval $[0, 1]$, we apply a transformation to map these values to $\mathbb{R}$, facilitating unconstrained regression. Among

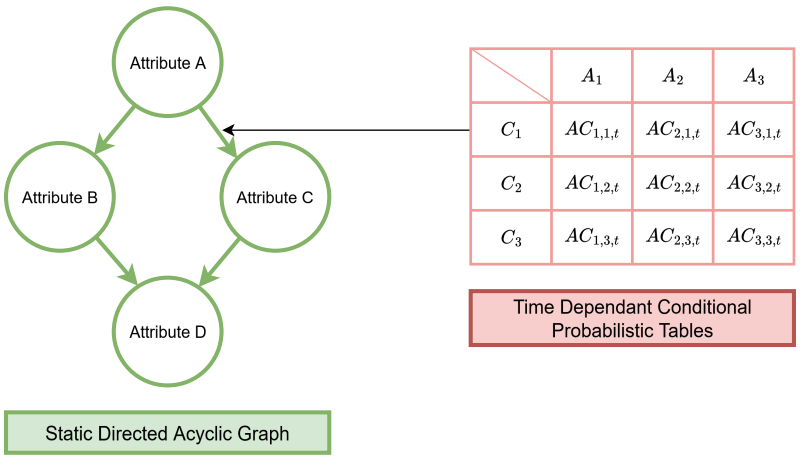


Figure 2: The two steps to learn a distribution with Bayesian Network: the Directed Acyclic Graph is static and the Conditional Probabilistic Tables are dynamic in our projection model.

various encodings, we choose the following encoding $\tilde{c}_{i,t}$:

$$\tilde{c}_{i,t} = \operatorname{arctanh}(c_{i,t} - c_{i,m}), \tag{1}$$

where $c_{i,m}$ corresponds to the most recent observed year y_m, serving as the reference baseline. This encoding captures the relative shift from the last observed state.
To project these values, a regression model $r_i(t)$ is fitted to the encoded coefficients $\tilde{c}_{i,t}$. The model r_i is specified either as a constant (order 0), corresponding to a weighted aggregation of past values, or as an affine (order 1) function, capturing a time-evolving trend. Restricting the regression to low-order polynomials is a deliberate choice to mitigate overfitting, as higher-degree functions are prone to capturing stochastic noise rather than meaningful structural dynamics. Consequently, this projection framework is best suited for short- to medium-term horizons, since linear trends cannot be reliably extrapolated over the long term without risking divergence from underlying demographic realities.
The regression parameters, i.e. the r_i coefficients, are estimated by minimizing a weighted least squares objective:

$$\mathcal{L}_{\alpha,n}(r_i) = \sum_{t=1}^{m} (1-\alpha)^{y_P - t} (r_i(t) - \tilde{c}_{i,t})^2; r_i \in \mathbb{R}_n[X], \tag{2}$$

where $\alpha \in [0,1]$ is a discount (or forgetting) factor, $n \in \{0,1\}$ denotes the polynomial degree, and y_P is the target year. The α parameter gives more weight to recent observations, ensuring that the projection is more influenced by recent trends than by historical fluctuations.
Finally, the projected coefficient for the target year y_P is recovered through the inverse transform:

$$c_{i,y_P} = \tanh(r_i(y_P)) + c_{i,m}. \tag{3}$$

Walk-Forward Sliding Window validation for hyperparameters selection in coefficient projections The accuracy of the projected coefficients depends on two critical hyperparameters introduced in Equation 2: α and n. These parameters are optimized through a Walk-Forward Sliding Window validation process, which assesses predictive performance while preserving both the temporal structure and the forecast horizon of the final projection.
As illustrated in Figure 3, the historical data are partitioned to maintain temporal symmetry between the validation and prediction phases. For example, in the case of a three-year forecast horizon targeting 2021, observations from 2011 to 2015 are used to calibrate the hyperparameters by validating against the year 2018. The resulting optimal parameters are then applied to the 2014–2018 interval to generate the projection for 2021.
Because optimal settings are variable-specific, hyperparameter selection is conducted independently for each CPT. Consequently, all coefficients within a single CPT share the same α and n. To ensure computational tractability, a grid search is performed sequentially across the tables (following the topological ordering from the DAG) rather than through a high-dimensional global joint optimization.

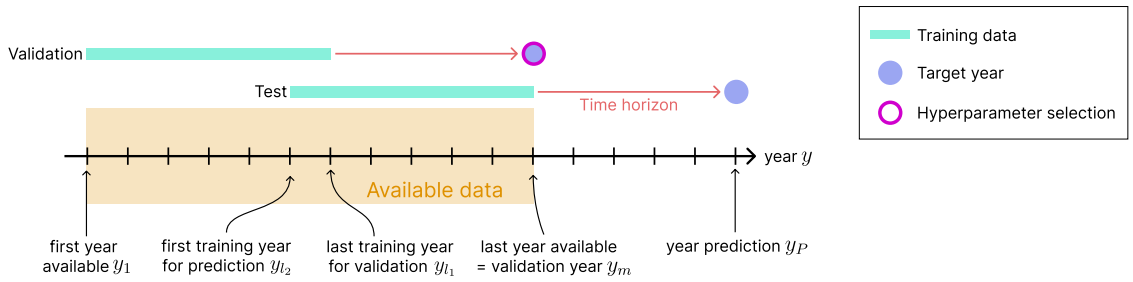


Figure 3: Data partitioning for validation and prediction. The training intervals and projection horizons are maintained symmetrically to ensure the validation phase accurately simulates the target prediction task.

In the CPT evaluation process, a direct parametric coefficients comparison is insufficient, as the impact of each coefficient on the final distribution depends on the relative frequency of the conditional states. Therefore, we align our selection process with the primary benchmark of population synthesis: the distributional accuracy of the resulting synthetic population. Furthermore, as the CPTs govern the relationships between attributes, the validation must prioritize multivariate consistency over the isolated marginals.
Our validation method evaluates this accuracy by leveraging the sequential nature of Bayesian Network sampling, following the same sequence to fit the hyperparameters. As attributes are generated hierarchically according to the DAG, we optimize the hyperparameters for a given table by comparing the exact joint distribution of the target attribute and the previously fixed attributes, against the corresponding joint distribution in the validation set.
Evaluation is performed using the average Standardized Root Mean Squared Error $\overline{SRMSE}$ (Lederrey et al., 2022; Darsel, Côme, & Oukhellou, 2026). For each table, we calculate the $\overline{SRMSE}_3$ (trivariate) to capture complex interactions; where the sub-graph contains fewer than three variables, the metric is adapted to $\overline{SRMSE}_2$ or $\overline{SRMSE}_1$ accordingly. Once the hyperparameters yielding the best distribution score have been identified for a table, they are fixed, and the process moves to the next CPT in topological order.

IPF postprocessing

While the projected Bayesian Network captures the evolving joint distribution, we add an additional refinement step to ensure alignment with marginals through a projective Iterative Proportional Fitting (IPF) procedure. IPF is a popular reweighting method to satisfy predefined marginal constraints (Deming & Stephan, 1940), which has seen extensive success in population synthesis (Beckman et al., 1996; Ye et al., 2009). In our model, we apply this technique by imposing projected marginals onto the synthetic agents previously sampled from the projected BN.
To compute these projected marginals, we adopt a methodological framework consistent with the projection of the CPTs, utilizing the same encoding, regression model, and hyperparameter tuning protocol. However, unlike the CPT validation process, the hyperparameters for marginals are selected by minimizing the Mean Squared Error (MSE) of the coefficients. This approach is appropriate in this context because the projected coefficients directly represent the final target frequencies, whereas CPT validation requires a multivariate metric to ensure consistency.

3 Experiments, Results & Discussion

Experimental Setup

To evaluate the robustness of our projection framework, we conducted a series of experiments across varying temporal horizons (one to five years), attribute dimensionalities (5 to 11 variables), and training sizes (0.03% and 1% of the total population). To ensure comparability across different projection scales, we utilize a five-year rolling window of training data for each projection model, independent of the validation phase.
We evaluate the distribution and realism of the generated population, using metrics proposed by Darsel, Côme, & Oukhellou (2026). The distribution is evaluated through the average Standard-

ized Rooted Mean Squared Error ($\overline{SRMSE}$) for univariate, bivariate and trivariate distributions. Regarding the realism, we consider the Share of Samples with a Couple of Instances Out of Data ($SSCIOD$), which identifies inconsistent attribute combinations.
Our model's performance is benchmarked against three reference cases:

- **Static BN:** A Bayesian Network trained solely on the most recent available year y_m (capturing the state but ignoring temporal trends).
- **Target-Year BN:** A BN trained directly on the target year's data y_P (representing the theoretical upper bound of the BN structure).
- **Raw Training Data:** A direct comparison with the most recent training dataset y_m to evaluate the baseline representativeness of a "no-model" approach for future years.

Data Sources and Partitioning

For our experiments, we use the open-source French Census data which is published annually. We consider the data from 2007 to 2021[2].
All projection scenarios are evaluated against the 2021 census as the target year. The allocation and role of the various yearly datasets are summarized in Table 1.

Time Horizon (y)	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
1									O	O/X	O/X	O/X	O/X	V/X	T
2							O	O/X	O/X	O/X	O/X	X	V/X		T
3					O	O	O	O/X	O/X	X	X	V/X			T
4			O	O	O	O	O/X	X	X	X	V/X				T
5	O	O	O	O	O	X	X	X	X	V/X					T

Table 1: Dataset roles across experimental scenarios (T: Testing, X: Training, V: Validation, O: Training for Validation)

Beyond temporal horizons, we evaluate the impact of data complexity using four incremental dimensionality cases:

- **5 attributes**: Age, Sex, Diploma, Household size and Department
- **7 attributes**: The previous set + Socioprofessional category and number of cars
- **9 attributes**: The previous set + Martial status and Accommodation type
- **11 attributes**: The previous set + Activity type and Household type

Results

The performance of our model is given in Figure 4.
Regarding the realism of the synthetic population, the $SSCIOD$ for the projected models remain largely consistent with the baseline BNs, with significant divergence in only two scenarios at the 5-year horizon.
The distributional accuracy results from Figure 4 have been aggregated in Figure 5, for each metric and training size. Each data point represents the performance of the projected model relative to the static baseline, normalized by the performance of the target-year BN. In this visualization, points falling below the identity line indicate a successful improvement over the static model. Our framework consistently improves marginal distributions across all scenarios. In multivariate distributions, the model outperforms the static baseline in 36 out of 40 experimental setups. We observe that the ratio scores are generally higher for the 1% training size, as the target-year reference achieves higher absolute performance at this scale. Notably, with a training size of 0.03%, the projection model occasionally outperforms both the raw training data and the target-year BN, demonstrating the advantage of incorporating multi-year data trends to mitigate sampling noise.
Although our models always improve the marginal distributions, we observe that only 4 experiences out of the 40 scenarios fail to improve multivariate distributions. From Figure 4, we observe that 2 of these experiments correspond to cases where the realism has exploded. When we compare the number of attributes in the SRMSE, we observe that the multivariate evaluation contracts the ratio scores towards 1. These ratios have higher values for the 1% training size because the baseline

[2]Details on the content and processing of the datasets are available here.

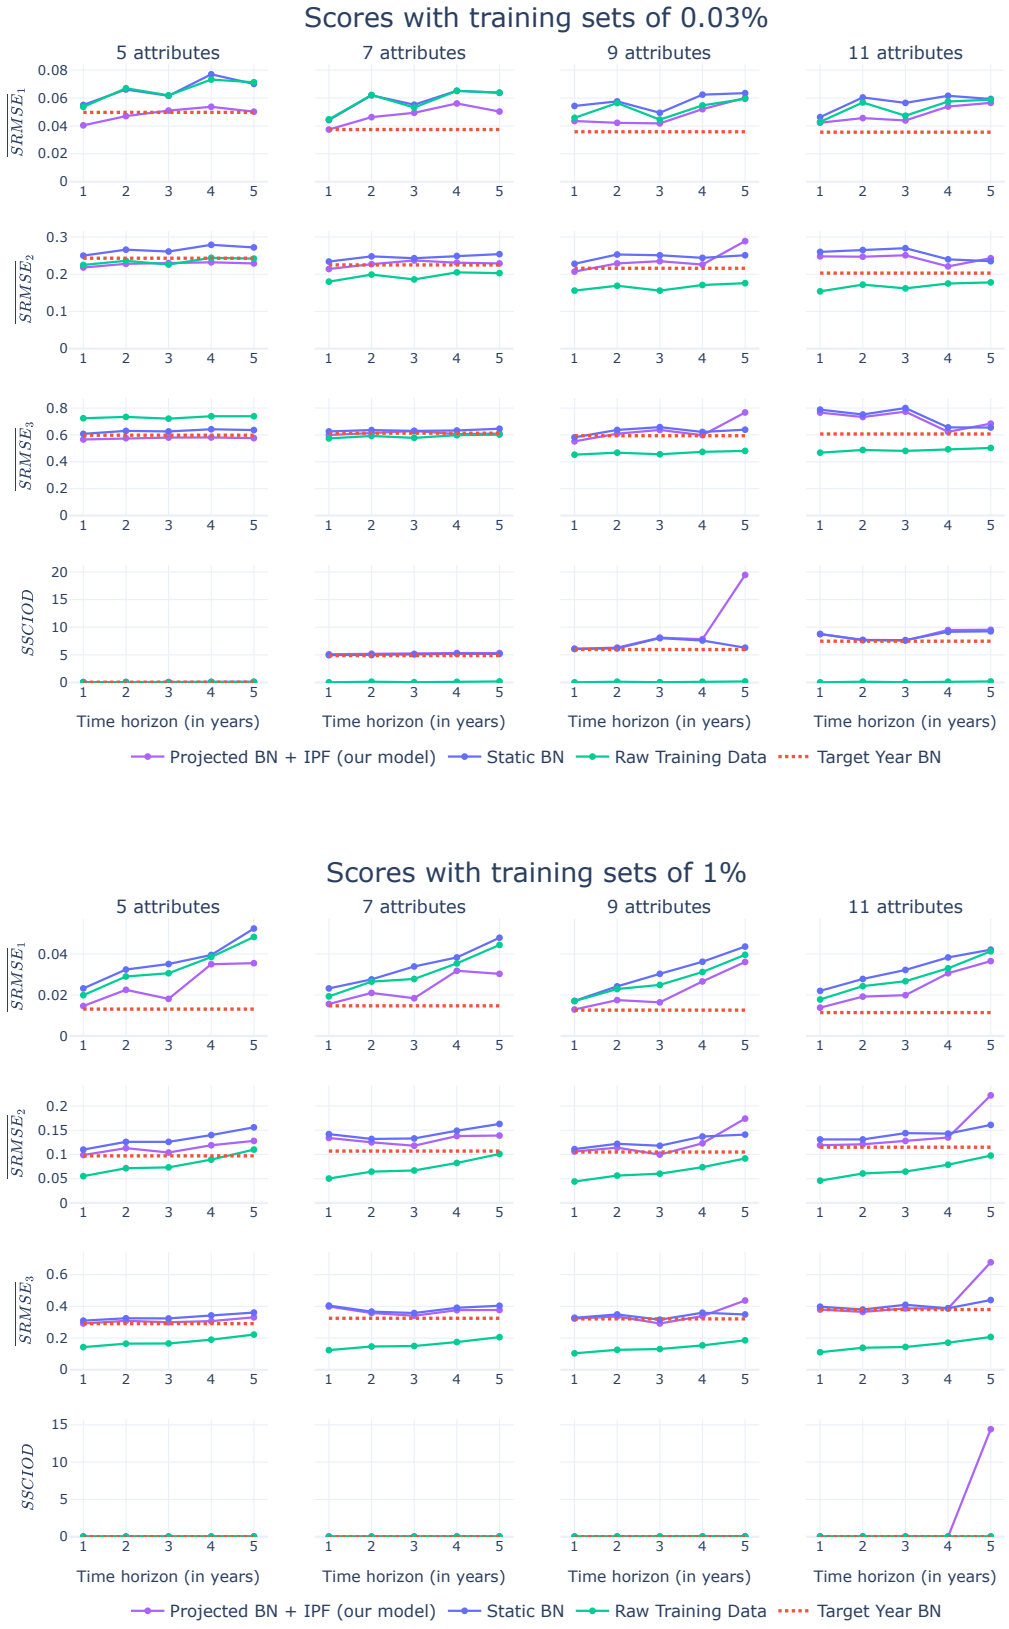


Figure 4: Comparison of the distribution and the realism of synthetic populations for different time horizons and synthetic population forecasting models

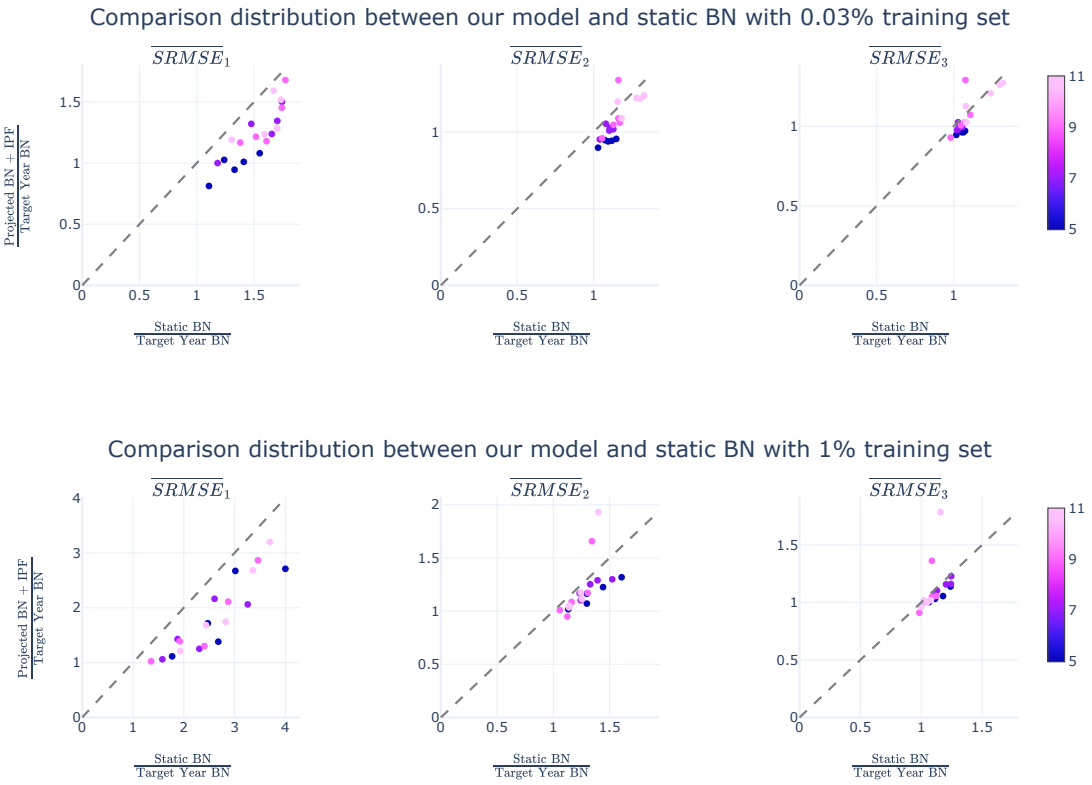


Figure 5: Comparison distribution scores between our model and static BN relatively to the scores from target-year BN depending on the number of attributes

and the target-year BN model achieve high performance. At 0.03%, the projection model surpasses the training data and the target-year BN model, demonstrating the value of combining years in the training process. Furthermore, we observe that multivariate evaluations tend to contract the ratio scores towards 1.

Discussions

The results suggest several key insights into the behaviour of temporal population projection. In Figure 5, we observe that as the number of attributes increases, the relative gain of the methodology tends to decrease. This likely represents a manifestation of the curse of dimensionality in Bayesian Networks, a phenomenon previously documented in population synthesis literature (Borysov et al., 2019), where the complexity of the joint distribution increases exponentially with the number of variables.

Two out of the four experiences that do not outperform the static baseline in terms of multivariate distribution also exhibit poor realism scores. As the realism of a sample is independent from the specific year, evaluating the realism of the generated population on the training data can serve as a reliable guardrail to detect unsuccessful projections. If the realism score ($SSCIOD$) of the projected population deviates significantly from the baseline BN, the projection should be considered unstable. This allows for an internal validation mechanism: rejecting projections where the learned temporal shifts lead to logically impossible attribute combinations. In the 2 experiments with low realism performances, the $SSCIOD$ of the model evaluated on the validation data is 39.2% (9 variables, 0.03%) and 14.7% (11 variables, 1%), confirming they would have been rejected by this validation process.

Finally, Tables 2 and 3 present the proportion of parameters projected using linear (order 1) versus constant (order 0) models for the BN and IPF steps, respectively. In nearly all scenarios, the optimal solution utilizes a hybrid of both, illustrating that the proposed validation procedure effectively captures attribute-dependent tendencies by identifying which distributions evolve linearly and those which can be estimated through a weighted averaging of observed frequencies. For the latter, accessing to multiple years of data allows to aggregate information and get a thinner

		Time horizon (in years)				
Training size	# Attributes	1	2	3	4	5
0.03%	5	20.0%	0.0%	40.0%	20.0%	0.0%
	7	28.57%	14.29%	28.57%	14.29%	0.0%
	9	33.33%	0.0%	0.0%	22.22%	11.11%
	11	27.27%	18.18%	27.27%	18.18%	27.27%
1%	5	40.0%	60.0%	40.0%	60.0%	60.0%
	7	57.14%	57.14%	71.43%	57.14%	71.43%
	9	44.44%	44.44%	55.56%	44.44%	66.67%
	11	63.64%	27.27%	54.55%	54.55%	54.55%

Table 2: Rate of linear projection for the conditional probability tables from the Bayesian Network

		Time horizon (in years)				
Training size	# Attributes	1	2	3	4	5
0.03%	5	40.0%	0.0%	40.0%	40.0%	20.0%
	7	42.86%	0.0%	42.86%	42.86%	14.29%
	9	55.56%	0.0%	55.56%	33.33%	11.11%
	11	45.45%	0.0%	45.45%	36.36%	9.09%
1%	5	80.0%	60.0%	60.0%	60.0%	40.0%
	7	71.43%	57.14%	71.43%	57.14%	57.14%
	9	66.67%	55.56%	66.67%	55.56%	55.56%
	11	72.73%	54.55%	63.64%	63.64%	54.55%

Table 3: Rate of linear projection for the Iterative Proportional Fitting coefficients.

coefficient estimation. For instance, at 0.03% training size, the presence of experiences with only constant projections suggests that the model defaults to a weighted average to filter out the high variance and stochastic noise inherent in small datasets. In such cases, linear regression may fail to capture a trend due to significant variations in the coefficients between successive points, whereas the order-0 model provides a more robust estimate.

4 Conclusions & Perspectives

In this paper, we introduce time-evolving variants of Bayesian Networks and the Iterative Proportional Fitting procedure. Combined, these components form a unified framework for synthetic population forecasting that leverages multi-year data and operates independently of external datasets. Attribute-level calibration enables the model to capture attribute-specific dynamics, distinguishing temporal evolution from stable relationships, for which multi-year data provides more robust static estimates. For validation, we employ a data partitioning strategy tailored to replicate projection constraints.

The results indicate that our model improves distributional accuracy over static baselines without sacrificing realism in most cases. Although performance decreases with longer forecasting horizons, the approach performs well for low-dimensional attribute sets and reduces variance in sparse data regimes. Evaluating realism against training data further allows the detection and rejection of unreliable projections. An ablation study was performed but is omitted due to space constraints and will be included in a future extended version.

Despite its strengths, the model relies on regression-based components and is therefore primarily suited to capturing trends. Over longer horizons, linearity may lead to unrealistic extrapolations, suggesting the need for alternative approaches, such as conditioning on macro-level indicators. Moreover, the model does not account for structural breaks in individual behavior caused by external shocks or changing mindsets. Extending it to more complex network structures could also better capture higher-order dependencies across attributes.

References

Beckman, R. J., Baggerly, K. A., & McKay, M. D. (1996, November). Creating synthetic baseline populations. *Transportation Research Part A: Policy and Practice*, *30*(6), 415–429. Retrieved 2024-04-05, from https://www.sciencedirect.com/science/article/pii/0965856496000043 doi: 10.1016/0965-8564(96)00004-3

Borysov, S. S., Rich, J., & Pereira, F. C. (2019, September). How to generate micro-agents? A deep generative modeling approach to population synthesis. *Transportation Research Part C: Emerging Technologies*, *106*, 73–97. Retrieved 2024-04-05, from https://www.sciencedirect.com/science/article/pii/S0968090X1831180X doi: 10.1016/j.trc.2019.07.006

Darsel, V., Côme, E., & Oukhellou, L. (2025, June). Population Synthesis with Deep Generative Models -is it worth it? Exploring new models and metrics. In *13th Symposium of the European Association for Research in Transportation.* Munich, Germany: Technische Universität München. Retrieved 2025-07-03, from https://hal.science/hal-05130079

Darsel, V., Côme, E., & Oukhellou, L. (2026, May). Robust and reproducible evaluation framework for population synthesis models—Application to probabilistic and deep generative models. *Transportation Research Part C: Emerging Technologies*, *186*, 105627. Retrieved 2026-03-19, from https://www.sciencedirect.com/science/article/pii/S0968090X26001154 doi: 10.1016/j.trc.2026.105627

Darsel, V., Côme, E., Rich, J., Pereira, F. C., & Oukhellou, L. (2026, July). The Fragility Of Synthetic Populations Over Time: An Evaluation Of Forecasting Strategies For Agent-Based Models. In *World Conference on Transport Research 2026 [Accepted].* Toulouse, France: Transportation Research Procedia.

Deming, W. E., & Stephan, F. F. (1940). On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known. *The Annals of Mathematical Statistics*, *11*(4), 427–444. Retrieved 2026-04-13, from https://www.jstor.org/stable/2235722

Fatmi, M. R., & Habib, M. A. (2017, January). Baseline Synthesis and Microsimulation of Life-stage Transitions within an Agent-based Integrated Urban Model. *Procedia Computer Science*, *109*, 608–615. Retrieved 2025-05-01, from https://www.sciencedirect.com/science/article/pii/S1877050917310359 doi: 10.1016/j.procs.2017.05.366

Heckerman, D. (1998, March). A Tutorial on Learning With Bayesian Networks.

Koller, D., & Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques.* MIT Press. (Google-Books-ID: 7dzpHCHzNQ4C)

Kukic, M., & Bierlaire, M. (2025, September). Adaptive synthetic generation using one-step Gibbs Sampler. *Transportation Research Interdisciplinary Perspectives*, *33*, 101597. Retrieved 2026-01-08, from https://www.sciencedirect.com/science/article/pii/S2590198225002763 doi: 10.1016/j.trip.2025.101597

Lederrey, G., Hillel, T., & Bierlaire, M. (2022, March). *DATGAN: Integrating expert knowledge into deep learning for synthetic tabular data.* arXiv. Retrieved 2025-01-23, from http://arxiv.org/abs/2203.03489 (arXiv:2203.03489 [cs]) doi: 10.48550/arXiv.2203.03489

Lomax, N., Smith, A. P., Archer, L., Ford, A., & Virgo, J. (2022). An Open-Source Model for Projecting Small Area Demographic and Land-Use Change. *Geographical Analysis*, *54*(3), 599–622. Retrieved 2025-04-29, from https://onlinelibrary.wiley.com/doi/abs/10.1111/gean.12320 (_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/gean.12320) doi: 10.1111/gean.12320

Namazi-Rad, M.-R., Mokhtarian, P., & Perez, P. (2014). Generating a Dynamic Synthetic Population – Using an Age-Structured Two-Sex Model for Household Dynamics. *PLOS ONE*, *9*(4), e94761. Retrieved 2025-07-01, from https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0094761 doi: 10.1371/journal.pone.0094761

Prédhumeau, M., & Manley, E. (2023, March). A synthetic population for agent-based modelling in Canada. *Scientific Data*, *10*(1), 148. Retrieved 2025-04-28, from https://www.nature.com/articles/s41597-023-02030-4 doi: 10.1038/s41597-023-02030-4

Rich, J. (2018, December). Large-scale spatial population synthesis for Denmark. *European Transport Research Review*, *10*(2), 63. Retrieved 2026-04-01, from https://doi.org/10.1186/s12544-018-0336-2 doi: 10.1186/s12544-018-0336-2

Sun, L., & Erath, A. (2015, December). A Bayesian network approach for population synthesis. *Transportation Research Part C: Emerging Technologies*, *61*, 49–62. Retrieved 2024-04-05, from https://www.sciencedirect.com/science/article/pii/S0968090X15003599 doi: 10.1016/j.trc.2015.10.010

Ye, X., Konduri, K., Pendyala, R., Sana, B., & Waddell, P. (2009, January). Methodology to match distributions of both household and person attributes in generation of synthetic populations. In *88th Annual Meeting Transportation Research Board.* Retrieved 2024-04-05, from https://trid.trb.org/View/881554