

Research Article

World Journal of Tourism Management

Predicting Night Cultural Tourism Consumption and Generating Personalized Tourism Scenes: An Evolutionary-Optimized Conditional GAN Framework

Oleksandr P. Krupskyi¹, Yuliia M. Stasiuk² and Maria Nascimento Cunha³

¹PhD in Psychology, Associate Professor at Oles Honchar Dnipro National University, Dnipro, Ukraine

²Oles Honchar Dnipro National University, Dnipro, Ukraine

³Ph.D. in Business Sciences, Professor at ISMT-Instituto Superior Miguel Torga, Coimbra, Portugal

***Corresponding Author**

Oleksandr P. Krupskyi, PhD in Psychology, Associate Professor at Oles Honchar Dnipro National University, Dnipro, Ukraine.

Submitted: 2026, Aug 07; Accepted: 2026, Sep 28; Published: 2026, Oct 11

Citation: Krupskyi, O. P., Stasiuk, Y. M., Cunha, M. N. (2026). Predicting Night Cultural Tourism Consumption and Generating Personalized Tourism Scenes: An Evolutionary-Optimized Conditional GAN Framework. *World J Tourism Mgm*, 2(3), 01-16

1. Introduction

The night-time economy has emerged as an increasingly important component of contemporary urban development, contributing to tourism activity, consumption opportunities, and urban vitality [1]. Within this broader economic ecosystem, night cultural tourism represents a distinctive domain of consumption, encompassing a wide range of activities and experiences, including sightseeing, cultural participation, gastronomy, shopping, entertainment, and leisure [2]. Recent research indicates that night-time tourism consumption is shaped by a combination of spatial, socioeconomic, infrastructural, and technological factors. Xiong et al for example, identified significant spatial disparities in national night cultural tourism consumption agglomeration zones in China, highlighting the influence of social development, tourism industry development, economic conditions, and transportation [1]. Similarly, Zhang et al. demonstrated the value of integrating multi-source data to assess the spatial distribution and suitability of nighttime tourism environments [2].

The increasing digitalization of tourism further complicates the dynamics of night-time consumption. Digital technologies can modify accessibility, information flows, mobility patterns, and the interaction between consumers and tourism environments. Recent evidence suggests that digital mobility platforms can reshape the spatial distribution of nighttime hospitality consumption. Xie et al, using large-scale ride-hailing data, found that digital mobility services can contribute to the spatial diffusion of nighttime hospitality consumption, while the temporal evolution of consumption follows a nonlinear pattern [3]. Such findings reinforce the view that night-time tourism is a dynamic

consumption ecosystem in which technological, spatial, temporal, and behavioural dimensions interact. More broadly, Liu et al. demonstrate that digital technology can influence the resilience and adaptive capacity of tourism systems, further illustrating the growing relevance of digital technologies in contemporary tourism development [4].

From a tourism marketing perspective, these developments highlight the need to move beyond descriptive approaches to night-time tourism and towards more sophisticated methods for understanding and anticipating consumer behaviour. Consumers' night-time tourism decisions emerge from the interaction of individual preferences, contextual conditions, temporal factors, destination characteristics, and experiential stimuli. The resulting behavioural data may therefore exhibit nonlinear relationships, temporal dependencies, heterogeneous structures, and high-dimensional characteristics. These properties create methodological challenges for conventional analytical approaches and increase the potential value of computational intelligence methods capable of identifying complex patterns and optimizing predictive models.

Evolutionary computation provides a relevant methodological foundation for addressing complex optimization problems. Recent advances have focused on adaptive mechanisms capable of balancing exploration and exploitation, maintaining solution diversity, and improving convergence in complex search spaces [5,6]. Chen et al, for instance, proposed a tri-population evolutionary algorithm incorporating adaptive stage-switching to dynamically balance diversity and convergence in constrained

multi-objective optimization [5]. Similarly, Silva et al. developed adaptive evolutionary strategies incorporating positive-adaptive mutation and demonstrated their applicability to complex optimization problems [6]. These developments suggest that evolutionary optimization can provide a useful complementary mechanism for feature selection and parameter optimization in predictive modelling. In the context of night cultural tourism, such mechanisms may facilitate the identification of relevant behavioural features and the optimization of predictive architectures designed to capture complex consumption patterns.

However, predicting consumer preferences represents only one component of the challenge faced by contemporary tourism destinations. The translation of behavioural insights into relevant and attractive experiences is equally important. In experience-oriented tourism environments, visual and spatial characteristics can contribute to how consumers perceive, interpret, and engage with destinations. Consequently, the capacity to generate alternative representations of tourism environments in response to identified consumer preferences may constitute an important extension of predictive tourism analytics. Recent developments in generative artificial intelligence have expanded the possibilities for automated visual content generation. Generative adversarial networks (GANs) have demonstrated considerable potential for image synthesis and the generation of realistic visual representations. Contemporary research has explored GAN architectures for high-fidelity and photorealistic image generation, image restoration under heterogeneous degradation conditions, and spatio-temporal modelling and forecasting [7-9].

Other studies have demonstrated the applicability of conditional GAN architectures to the generation of data controlled by specific characteristics or conditions [10]. More recently, hybrid GAN architectures have sought to improve stability, feature representation, and robustness, while attention-based GAN architectures have further advanced the generation of photorealistic images [11,12]. Collectively, these developments indicate that generative models can provide flexible mechanisms for producing synthetic content conditioned by relevant input characteristics. Despite these advances, the predictive and generative dimensions of tourism consumption remain conceptually and technologically fragmented. Research on nighttime tourism has increasingly employed spatial analysis, multi-source data, and digital technologies to understand consumption environments, while computational research has developed increasingly sophisticated evolutionary optimization and generative models [[1-12]. Nevertheless, comparatively limited attention has been given to integrating these approaches into a consumer-oriented framework capable of linking behavioural prediction with the generation of personalized tourism scenarios.

To address this gap, the present study proposes an integrated computational framework combining an Alpha Evolution-optimized consumer behaviour prediction model with a conditional generative adversarial network (cGAN). The predictive component is designed to extract relevant patterns from historical consumer

behaviour data and estimate potential preferences and demand characteristics associated with night cultural tourism. These predicted characteristics are subsequently introduced as conditional information into the generative component, enabling the generation of visual representations of night cultural tourism environments aligned with the identified consumer profiles and preferences. The proposed framework therefore establishes a conceptual and technological connection between two traditionally distinct stages of tourism marketing: **consumer intelligence and experience design**. Rather than treating prediction and content generation as independent processes, the framework conceptualizes them as interconnected components of a data-driven personalization process.

Historical behavioural data inform prediction; predicted consumer characteristics guide generative processes; and the resulting tourism scenarios provide a basis for designing more contextually relevant experiences. In this sense, the study contributes to the emerging intersection of tourism marketing, consumer analytics, evolutionary computation, and generative artificial intelligence. By integrating behavioural prediction with conditional visual generation, the proposed framework seeks to extend the application of artificial intelligence in tourism from analytical and predictive functions towards the proactive design of consumer-oriented experiences. Its contribution lies not only in the application of individual computational techniques, but in their integration into a unified architecture designed to connect **what consumers are likely to prefer with how tourism environments can be represented and designed in response to those preferences**. This perspective provides a foundation for further research into AI-enabled personalization and the digital transformation of night cultural tourism marketing.

1.1. Overview of Generative Adversarial Network Principles

Generative adversarial networks (GANs) are deep generative models designed to learn the underlying distribution of observed data and generate new synthetic samples that reproduce relevant characteristics of the original data [13–15]. A conventional GAN comprises two neural networks with complementary functions: a generator and a discriminator [16]. The generator, (G), transforms a latent noise vector (z) into a synthetic sample (G(z)), whereas the discriminator, (D), evaluates whether an input sample originates from the real data distribution or has been generated by (G) [17]. Through adversarial training, the two networks are optimized iteratively, enabling the generator to progressively approximate the target data distribution [18]. The learning process can be formulated as a two-player minimax optimization problem [19]. Given a latent vector (z) sampled from a prior distribution ($p_z(z)$), the generator produces a synthetic observation (G(z)). The discriminator receives either a real observation (x), sampled from the empirical distribution ($p_{data}(x)$), or a generated observation (G(z)), and estimates the probability that the input belongs to the real data distribution [20]. The original GAN objective is therefore expressed as:

$$\mathbb{E} \{x \sim p_{data}(x)\} [\log D(x)]$$

$$+ \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))].$$

where $(p_{data}(x))$ denotes the distribution of real observations and $(p_z(z))$ represents the prior distribution of the latent variable (z) .

During training, the generator and discriminator are updated alternately. For a fixed generator (G) , the optimal discriminator is given by:

$$\frac{p_{data}(x)}{p_{data}(x) + p_g(x)},$$

where $(p_g(x))$ denotes the distribution induced by the generator. At the theoretical equilibrium of the adversarial optimization process, the generated distribution converges to the real data distribution, such that $(p_g(x) = p_{data}(x))$, and the discriminator assigns a probability of 0.5 to observations originating from either distribution.

For applications in which the generated output must reflect specific characteristics, the conventional GAN can be extended into a conditional generative adversarial network (cGAN). In a cGAN, an additional conditional variable (y) is incorporated into both the generator and discriminator, allowing the generation process to be explicitly guided by predefined information [16,17]. The corresponding objective function can be expressed as:

$$\mathbb{E}_{x \sim p_{data}(x)} [\log D(x|y)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z|y)))].$$

Here, (y) represents the conditional information used to guide the generation process. Depending on the application, (y) may include categorical attributes, textual descriptors, behavioural characteristics, or other contextual variables. This conditional mechanism is particularly relevant to the present study because it provides a means of incorporating predicted consumer characteristics into the generation of tourism-related visual content. Rather than generating night cultural tourism scenes independently of consumer information, the cGAN can be conditioned on specific preference profiles identified by the predictive component.

An important development of the original GAN architecture is the Wasserstein generative adversarial network (WGAN), which employs the Wasserstein distance to quantify the discrepancy between the real and generated data distributions [21–23]. Through the Kantorovich–Rubinstein duality, the Wasserstein distance can be formulated as:

$$\mathbb{E}_{z \sim p_z(z)} [f(G(z))]$$

where $(\|f\|_L \leq 1)$ represents the Lipschitz constraint imposed on the critic function (f) . By providing a more informative measure of the discrepancy between probability distributions, the Wasserstein formulation can contribute to more stable adversarial optimization and more informative gradient estimation [21–23]. The flexibility of GAN architectures has supported their application to a broad range of generative tasks, including image synthesis, image restoration, spatio-temporal prediction, and conditional data generation [24–26]. Recent developments have further incorporated attention mechanisms, hybrid architectures, and additional conditioning strategies to improve the quality, stability, and controllability of generated outputs. These developments are particularly relevant to tourism marketing because they enable generative systems to move beyond the production of generic visual content towards the generation of content conditioned on specific consumer or contextual characteristics.

In the context of the present research, the principal relevance of GAN technology therefore lies in its capacity to establish a computational link between **consumer information and tourism experience representation**. When consumer preferences are introduced as conditional variables, the generative component can be used to produce visual representations of night cultural tourism environments aligned with specific consumer profiles. This provides the methodological basis for moving from consumer behaviour prediction towards the design of more personalized tourism experiences.

1.2. Analysis of the Alpha Evolution Algorithm

Evolutionary algorithms are population-based optimization techniques inspired by mechanisms of natural selection and evolutionary adaptation [27]. Rather than searching directly for a single optimal solution, these algorithms maintain a population of candidate solutions and iteratively improve their fitness through evolutionary operators such as selection, recombination, and mutation [28]. Their population-based search strategy makes them suitable for complex optimization problems involving nonlinear relationships, multiple local optima, and high-dimensional search spaces [29]. The general evolutionary process begins with the initialization of a population of candidate solutions. Each individual is evaluated according to a predefined fitness function, after which individuals are selected and transformed through evolutionary operators to generate a new population. This process is repeated iteratively until a predefined termination criterion is reached. The fundamental objective is to achieve an appropriate balance between **exploration**, understood as the search across different regions of the solution space, and **exploitation**, referring to the refinement of promising solutions. Alpha Evolution constitutes an evolutionary optimization approach in which this balance is supported by explicitly distinguishing between high-quality elite solutions and the remaining members of the population.

The population is divided into two groups: an **Alpha group**, comprising the (K) individuals with the highest fitness values, and a **Beta group**, containing the remaining (N-K) individuals. The Alpha group provides directional guidance to the search process, whereas the Beta group contributes to exploration and population diversity [30].

Following fitness evaluation and ranking, individuals in the two groups participate in the evolutionary search process. For each individual (β_i) belonging to the Beta group, a trial vector is generated through a directional evolutionary operation:

$$\beta_i + F(\alpha_k - \beta_i) + F(\beta_i - \beta_j), \quad \text{tag}\{5\}$$

where (α_k) represents an individual selected from the Alpha group, (β_j) denotes another individual from the Beta group, and (F) is the scaling factor controlling the magnitude of the evolutionary update.

The directional operation allows the search to be informed by high-fitness solutions while retaining information from the broader population. The resulting trial vector can subsequently undergo probabilistic mutation, introducing stochastic perturbations that increase population diversity and reduce the likelihood of premature convergence. For real-valued representations, Gaussian perturbation can be applied to individual dimensions of the trial vector:

$$v_{\{i,j\}} + \text{randn}(0, \sigma^2), \quad \text{tag}\{6\}$$

where (j) denotes the dimension of the solution vector and (σ^2) controls the magnitude of the mutation. The mutation probability (p_m) determines the probability that a given dimension will undergo perturbation. Following the evolutionary and mutation operations, an elitist retention mechanism is used to preserve high-quality solutions while incorporating newly generated candidates into the subsequent population. The iterative process therefore combines directional exploitation of promising regions with stochastic exploration of alternative regions of the search space. The relevance of this mechanism to the present study lies in the optimization of the consumer behaviour prediction component. Night cultural tourism consumption may involve multiple interacting variables, nonlinear relationships, and heterogeneous behavioural patterns. An evolutionary optimization strategy can consequently be used to search for effective configurations of model parameters and/or relevant behavioural features, depending

on the specific predictive architecture adopted. In this sense, Alpha Evolution functions as an optimization mechanism rather than as an independent consumer behaviour model.

The integration of Alpha Evolution with the predictive component also establishes the first stage of the proposed personalization framework. Historical consumer data are used to identify behavioural patterns and relevant predictive information; the evolutionary algorithm supports the optimization of the predictive model; and the resulting model generates estimates of consumer preferences or demand characteristics. These outputs can subsequently be transferred to the generative component as conditional information.

1.3. Integration of Alpha Evolution and CGAN for Night Cultural Tourism Personalization

The combination of evolutionary optimization and conditional generative modelling provides the conceptual foundation for the proposed framework. The two techniques perform complementary functions within the overall architecture. Alpha Evolution is used to support the optimization of the predictive component, whereas the cGAN is responsible for transforming consumer-related conditional information into generated visual representations.

The proposed process can therefore be conceptualized as a sequential pipeline:

Historical consumer data → behavioural feature identification → Alpha Evolution optimization → consumer preference prediction → conditional information → CGAN → night cultural tourism scene generation.

The predictive stage focuses on extracting meaningful patterns from historical consumption data and estimating consumer preferences or demand characteristics. The optimization stage seeks to identify an effective configuration of the predictive model and/or relevant behavioural features. The generative stage subsequently incorporates the resulting consumer information as conditional input, enabling the production of visual tourism scenarios associated with the predicted characteristics. From a tourism marketing perspective, this integration is particularly relevant because it connects two traditionally distinct activities: **consumer intelligence and experience design**. Consumer analytics provide information about what consumers are likely to value, while generative technologies provide a mechanism for translating this information into alternative representations of tourism experiences.

The proposed architecture therefore moves beyond the isolated prediction of consumer behaviour by considering how predictive insights can inform the design and visualization of consumer-oriented tourism environments. The framework should not, however, be interpreted as assuming that algorithmic prediction automatically produces superior tourism experiences. Rather, it provides a technological mechanism through which predicted consumer characteristics can inform the generation of alternative scenarios. The effectiveness of this mechanism ultimately

depends on the predictive validity of the model, the quality and representativeness of the underlying data, the fidelity and relevance of the generated content, and the extent to which consumers perceive the resulting experiences as relevant and valuable.

Accordingly, the proposed framework contributes to the emerging intersection of **tourism marketing, consumer analytics, evolutionary optimization, and generative artificial intelligence**. Its central proposition is that the combination of behavioural prediction and conditional content generation can create a more integrated approach to the personalization of night cultural tourism experiences. Rather than treating consumer analytics and generative AI as independent technological applications, the framework connects them within a common consumer-oriented architecture in which behavioural information becomes an input for experience-oriented content generation.

1.4. Alpha-Based Evolution And Gan Framework

As illustrated in Figure 1, the proposed framework comprises two interconnected modules: a **consumer behaviour prediction module** and a **night cultural tourism scene generation module**. The two modules are sequentially connected through a conditional data interface, establishing a computational pipeline from behavioural data analysis to personalized visual content generation. The central principle of the framework is that information extracted from consumer behaviour prediction is subsequently used as conditional information to guide the generation of tourism scene content. The consumer behaviour prediction module is based on a multilayer perceptron (MLP). Rather than relying exclusively on conventional gradient-based parameter optimization, the proposed framework incorporates the Alpha Evolution algorithm to optimize the parameter configuration of the predictive model.

Each individual in the evolutionary population represents a candidate parameter configuration of the MLP. For each candidate solution, the predictive model is evaluated on the validation data, and the resulting prediction loss is used as the fitness criterion. High-performing solutions are retained as elite individuals, while evolutionary operations are used to generate alternative parameter configurations. Directional crossover promotes the exploration of promising regions of the parameter space, whereas probabilistic mutation introduces stochastic variation and maintains population diversity. The evolutionary process is repeated until the predefined convergence or stopping criterion is reached. The scene generation module is implemented using a conditional generative adversarial network (cGAN). The generator receives two inputs: a latent noise vector (z), sampled from a standard normal distribution, and a condition vector (c) generated by the consumer behaviour prediction module. These inputs are jointly processed through the generator network, which progressively transforms the latent representation into a synthetic night cultural tourism scene. The discriminator receives either a real scene image or a generated image together with the corresponding condition vector. Its role is therefore twofold: to distinguish real images from generated images and to assess whether the generated content is consistent with the specified conditional information.

These inputs are jointly processed through the generator network, which progressively transforms the latent representation into a synthetic night cultural tourism scene. The discriminator receives either a real scene image or a generated image together with the corresponding condition vector. Its role is therefore twofold: to distinguish real images from generated images and to assess whether the generated content is consistent with the specified conditional information.

The overall architecture can consequently be represented as:
Multi-source consumer data → **MLP** → **Alpha Evolution optimization** → **predicted consumer characteristics** → **conditional vector (c)** → **cGAN** → **generated night cultural tourism scene**.




Figure 1: Two Interconnected Modules: a Consumer Behaviour Prediction Module And A Night Cultural Tourism Scene Generation Module

The architecture establishes a direct connection between **consumer intelligence and experience-oriented content generation**. The predictive module identifies relevant behavioural characteristics, while the generative module translates these characteristics into visual representations of tourism environments. The framework therefore operationalizes personalization as a sequential process in which predicted consumer preferences provide the informational basis for generating conditionally relevant tourism content.

2. Analysis of The Training Process

The training procedure comprises two sequential stages corresponding to the consumer behaviour prediction module and the scene generation module. Although the modules are optimized separately, they are connected through the conditional representation generated by the predictive component.

2.1. Consumer Behaviour Prediction

The first stage involves training the MLP-based consumer behaviour prediction model. The Alpha Evolution algorithm is used to optimize the model parameter configuration. An initial population of ($N=50$) candidate solutions is generated, with each individual (θ_i) representing a complete candidate configuration of the predictive model.

For each candidate solution, the MLP is evaluated on the validation set. Given a multiclass classification task, categorical cross-entropy is used as the optimization objective:

$$-\frac{1}{N} \sum_{j=1}^N \sum_{c=1}^C y_{j,c} \log(\hat{y}_{j,c}),$$

where (N) represents the number of observations in the evaluation set, (C) denotes the number of classes, ($y_{j,c}$) is the observed class indicator for observation (j), and ($\hat{y}_{j,c}$) represents the probability assigned by the model to class (c).

Because lower cross-entropy corresponds to better predictive performance, the evolutionary optimization process seeks parameter configurations that minimize ($F(\theta)$). Candidate solutions are ranked according to their validation loss, and the highest-performing configurations are retained as elite solutions. Directional evolutionary operations and probabilistic mutation are subsequently applied to generate new candidate configurations. This process continues until the predefined maximum number of evolutionary generations is reached or the convergence criterion is satisfied.

The final optimized MLP is then used to generate the consumer-related predictions that constitute the conditional information for the subsequent scene-generation stage.

2.2. Conditional Scene Generation

The second stage is based on a conditional GAN architecture. The generator (G) receives a latent vector (z) and the condition vector (c) obtained from the consumer behaviour prediction module:

$$\hat{x} = G(z|c),$$

where ($\hat{x}$) represents the generated tourism scene.

The discriminator (D) receives either a real scene (x) or a generated scene ($\hat{x}$), together with the corresponding condition vector (c). In order to improve the stability of adversarial training, the proposed implementation employs a Wasserstein objective with gradient penalty (WGAN-GP).

The discriminator/critic objective is defined as:

$$\mathbb{E}_{x \sim P_r} [D(x|c)] + \lambda \mathbb{E}_{\tilde{x}} \left[\left(\|\nabla_{\tilde{x}} D(\tilde{x}|c)\|_2 - 1 \right)^2 \right],$$

where (P_r) and (P_g) denote the real and generated data distributions, respectively, ($\tilde{x}$) represents a random interpolation between real and generated samples, and (λ) is the gradient-penalty coefficient. In the present implementation, ($\lambda=10$).

The generator is trained to maximize the critic's output for generated samples, or equivalently to minimize:

$$L_G = -\mathbb{E}_{z \sim P_z} [D(G(z|c)|c)].$$

The generator and discriminator are updated alternately during training. To further encourage semantic consistency between the generated scene and the conditional representation, an additional condition-consistency loss is incorporated:

$$c - \phi(G(z|c))^2,$$

where ($\phi(\cdot)$) denotes a pretrained feature-extraction network that maps the generated scene into the corresponding feature space.

The overall generator objective can therefore be represented as a combination of the adversarial and conditional-consistency objectives:

$$[L_{\{G, total\}} = L_G + \lambda_{\{cond\}} L_{\{cond\}},]$$

where ($\lambda_{\{cond\}}$) controls the contribution of the condition-consistency term. The adversarial training process continues until the quality, diversity, and conditional consistency of the generated scenes reach a stable level according to the predefined evaluation criteria.

3. Protocol

3.1. Dataset and Data Preprocessing

The experimental dataset integrates multiple heterogeneous data sources associated with night cultural tourism consumption. These sources include night-light remote-sensing information, points of interest (POI), online user check-in data, and historical consumption records. The integration of these sources enables the behavioural characteristics of night-time tourism consumption to be represented across spatial, temporal, infrastructural, and consumption-related dimensions. The datasets are first harmonized spatially and temporally. Spatial observations are aggregated into a common geographical grid, while temporal information is aligned according to a common timestamp structure. Continuous variables are subsequently standardized using Z-score normalization.

Categorical variables are transformed using one-hot encoding, while textual labels are converted into numerical representations through vectorization. The resulting dataset is divided chronologically into training, validation, and test subsets using a 60:20:20 ratio. The chronological split is adopted to reduce temporal leakage and to ensure that the predictive model is evaluated on observations that occur after the period used for model training. Where class imbalance is present, SMOTE is applied exclusively to the training data. This restriction is important because synthetic oversampling of validation or test observations could introduce information leakage and compromise the validity of the evaluation procedure.

3.2. Evaluation Metrics

The predictive component is evaluated using **accuracy and F1 score**. Accuracy measures the proportion of correctly classified observations across the test set, whereas the F1 score represents the harmonic mean of precision and recall and is particularly informative when class distributions are imbalanced. For the scene-generation component, **Inception Score (IS) and Fréchet Inception Distance (FID)** are employed. IS provides an indirect

assessment of the confidence and diversity of generated images based on a pretrained image-classification model, whereas FID measures the distance between the feature distributions of real and generated images. Lower FID values indicate greater similarity between the corresponding feature distributions. All performance metrics are calculated using independent test data that are not used during model optimization.

3.3. Experimental Settings

The consumer behaviour prediction module employs an Alpha Evolution population of 50 individuals and a maximum of 200 evolutionary generations. The crossover rate is set to 0.85, while the mutation rate is set to 0.02. The underlying predictive model is a three-layer fully connected neural network with hidden-layer dimensions of 256, 128, and 64 neurons, respectively. The MLP is initially trained using the Adam optimizer with a learning rate of 0.001, while the Alpha Evolution algorithm is responsible for the higher-level optimization of the candidate parameter configurations. The scene-generation module uses a conditional GAN architecture. Both the generator and discriminator are optimized using Adam with a learning rate of 0.0002 and a batch size of 32. The latent vector has a dimensionality of 100, while the condition embedding has a dimensionality of 64. The generator uses a fourfold progressive upsampling structure, and the discriminator employs Leaky ReLU activation with a negative slope of 0.2. The total number of training iterations is set to 50,000.

4. Experimental Results

4.1. Ablation Analysis

The results reported in Table 1 compare different configurations of the proposed framework. Experiment 1, which incorporates Alpha Evolution, conditional GAN generation, and the integrated framework, achieves an accuracy of 90.5%, an F1 score of 0.89, and an Inception Score of 9.2.

Experiment 2 removes the Alpha Evolution optimization component, while Experiment 3 removes the conditional generation component. Experiment 4 evaluates the framework without the integrated end-to-end configuration. The observed differences across these experimental configurations indicate that the removal of individual components is associated with changes in predictive and generative performance.

However, these results should be interpreted as an **ablation analysis rather than as definitive evidence that each component independently causes the observed performance differences**. Formal statistical testing and repeated runs would provide stronger evidence regarding the robustness of these differences.

Experiment	Alpha Evolution Optimization	Conditional GANs	End-to-End Framework	Prediction Accuracy	F1 Score	Inception Score
1	✓	✓	✓	90.50%	0.89	9.2
2	×	✓	✓	87.30%	0.85	8.5
3	✓	×	✓	89.00%	0.87	7.8
4	✓	✓	×	88.50%	0.86	8.9

Note. ✓ = component included; × = component excluded. Higher values indicate better performance in terms of prediction accuracy, F1 score, and Inception Score.

Table 1: Results of the Ablation Analysis of the Proposed Framework

4.2. Comparative Model Evaluation

Table 2 compares the proposed Alpha-GAN framework with alternative modelling approaches. The Alpha-GAN configuration achieves an accuracy of 90.5%, an F1 score of 0.89, an Inception Score of 9.2, and an FID score of 21.8. The traditional time-series model records an accuracy of 85.0% and an F1 score of 0.80. The remaining GAN-based configurations produce intermediate results

across the reported metrics. These findings indicate that, within the experimental conditions adopted in this study, the integrated Alpha-GAN framework is associated with higher predictive and generative evaluation scores than the alternative configurations examined. Nevertheless, claims of general superiority should be restricted to the models, dataset, parameter settings, and evaluation protocol used in the present experiment.

Experiment	Model Type	Prediction Accuracy	F1 Score	Inception Score	FID Score
1	Alpha-GAN	90.50%	0.89	9.2	21.8
2	Traditional Time-Series Model	85.00%	0.80	—	—
3	Basic GAN	88.00%	0.84	7.5	24.0
4	Enhanced GAN	89.50%	0.88	8.8	23.0
5	Standard GAN	87.00%	0.83	8.0	25.0

Note. FID = Fréchet Inception Distance; GAN = generative adversarial network. An em dash (—) indicates that the metric is not applicable or was not calculated for the corresponding model.

Table 2: Performance Comparison of the Proposed Model With Benchmark Models

4.3. Spatial and Temporal Consumption Patterns

Figure 2 presents the distribution of night cultural tourism consumption across different geographical contexts and time periods. Consumption levels differ across areas, with the highest observed values reported for shopping areas (92%) and cultural areas (89%). Shopping-related consumption represents a substantial proportion of activity in both urban centres and shopping areas, accounting for 69% and 81%, respectively. The

temporal distribution further indicates that food and beverage consumption reaches its highest observed level between 19:00 and 20:00, at approximately 37.4%, before declining to approximately 10.3% between 01:00 and 02:00. These temporal differences illustrate the heterogeneous nature of night-time consumption and provide contextual information for the behavioural prediction component.




Figure 2: Spatial and Temporal Distribution of Night Tourism Consumption

4.4. Behavioural Prediction Performance

Figure 3 compares the predictions generated by the four evaluated models against the observed night tourism consumption index. The Alpha-GAN model shows closer correspondence with the observed consumption patterns in the reported scenarios, including weekend peaks and special events.

For example, the observed weekend peak reaches 119.1, while the model produces a corresponding prediction close to the observed

pattern. Similarly, the reported holiday and event values indicate relatively small deviations between the Alpha-GAN predictions and the observed data.

These results suggest that the integrated model can capture temporal variation in night cultural tourism consumption under the experimental conditions considered.




Figure 3: Forecast Analysis of Night Tourism Consumption Index Model

Figure 4 further examines prediction errors across consumption categories and modelling approaches. Special-event consumption presents the highest average prediction error among the reported categories, at approximately 15, followed by accommodation, at

approximately 10. When the models are compared, the Alpha-GAN configuration presents the lowest reported average prediction error, approximately 2, whereas the traditional model presents an average error of approximately 12.




Figure 4: Analysis Of Various Consumption Forecast Errors By Different Models

4.5. Spatial Distribution and Consumer Mobility

Figure 5 illustrates the spatial distribution of high- and low-consumption areas and the relationship between residential and commercial-cultural areas. High-consumption hotspots are primarily concentrated around the central region represented by the coordinates approximately (400, 400), with an estimated radius of 200 m. Commercial and cultural areas are concentrated around

approximately (600, 400), with an estimated radius of 300 m.

These spatial patterns provide contextual information for understanding the relationship between location characteristics and night-time consumption. However, spatial concentration should not by itself be interpreted as evidence of a causal relationship between destination characteristics and consumer behaviour.




Figure 5: Classification Diagram of Urban Tourism Areas

Figure 6 presents movement trajectories for different consumer segments, including families, young tourists, business travellers, local residents, and cultural enthusiasts. The reported stay ratios differ across destination environments, including urban centres, cultural districts, historical sites, shopping districts, entertainment

areas, and waterfront locations. The figure also compares observed trajectories with predictions generated by the traditional model and the Alpha-GAN model. According to the reported results, the Alpha-GAN predictions show closer correspondence with the observed trajectories.




Figure 6: Analysis of Travel Trajectory and Prediction Model of Consumer Groups

Figure 7 examines prediction errors across different time periods and expenditure categories. The Alpha-GAN configuration presents the lowest reported prediction error, particularly during the 22:00–00:00 period, where the average error is approximately 5.

The expenditure analysis indicates that accommodation represents the highest average expenditure category, at approximately USD 300, whereas food and beverage expenditure represents the lowest, at approximately USD 100.




Figure 7: Analysis of Travel Trajectory and Prediction Model Of Consumer Groups

Figure 8 presents the spatial distribution and direction of population movement. The density map indicates a concentration of population activity in the central area, with the highest reported density reaching approximately 0.7. The flow-direction visualization further identifies areas characterized by concentrated movement.

Together, these spatial representations provide contextual information regarding population density and mobility patterns and may contribute to the interpretation of night-time tourism consumption dynamics. In the proposed framework, such information can be incorporated into the broader behavioural representation used for prediction and subsequent scenario generation.




Figure 8: Urban Crowd Density and Flow Analysis

5. Results

The experimental results indicate that the proposed Alpha Evolution-cGAN framework provides competitive performance across both consumer behaviour prediction and conditional night-scene generation. The framework achieved a prediction accuracy of 90.50% and an F1 score of 0.89, while the generated scenes obtained an Inception Score (IS) of 9.2 and a Fréchet Inception Distance (FID) of 21.8. Taken together, these results indicate that the framework can address both the predictive and generative components of night cultural tourism analysis within an integrated architecture. The ablation analysis provides further evidence regarding the contribution of the individual components of the framework. When Alpha Evolution optimization was removed, prediction accuracy decreased from 90.50% to 87.30%, while the F1 score decreased from 0.89 to 0.85.

The Inception Score also declined from 9.2 to 8.5. When the conditional GAN component was removed, prediction accuracy decreased to 89.00%, the F1 score to 0.87, and the Inception Score to 7.8. Similarly, removing the end-to-end configuration resulted in an accuracy of 88.50%, an F1 score of 0.86, and an Inception Score of 8.9. These differences indicate that the removal of individual components is associated with reductions in predictive and/or generative performance, supporting their inclusion in the integrated framework. The comparison with benchmark models provides additional evidence of the framework's performance under the evaluated experimental conditions. The Alpha Evolution-cGAN framework achieved the highest prediction accuracy (90.50%) and F1 score (0.89) among the evaluated models. The traditional time-series model achieved an accuracy of 85.00% and an F1 score of 0.80, whereas the Basic GAN, Enhanced GAN, and Standard GAN achieved accuracies of 88.00%, 89.50%, and 87.00%, respectively.

For generative performance, the proposed framework achieved an IS of 9.2 and an FID of 21.8. The Basic GAN, Enhanced GAN, and Standard GAN obtained IS values of 7.5, 8.8, and 8.0 and FID values of 24.0, 23.0, and 25.0, respectively. Since higher IS values and lower FID values generally indicate more favourable generative performance, these results suggest that the proposed framework achieved a comparatively favourable balance between image diversity and similarity to the reference distribution.

The spatial and temporal analyses further demonstrate substantial heterogeneity in night cultural tourism consumption. Higher consumption levels were concentrated in shopping and cultural areas, while food and beverage activity displayed a pronounced evening peak followed by a decline during the early morning period. The spatial analysis similarly identified concentrated consumption hotspots in urban commercial and cultural areas. These findings indicate that night cultural tourism consumption varies considerably across both geographical and temporal dimensions, supporting the relevance of analytical approaches capable of representing nonlinear and dynamic patterns.

Overall, the results provide empirical support for the feasibility of integrating evolutionary optimization with conditional generative modelling in night cultural tourism. However, the reported

performance differences should be interpreted within the scope of the experimental design. In particular, differences between models should not be described as statistically significant unless supported by repeated experiments and appropriate statistical testing.

6. Discussion

The findings contribute to the intersection of consumer analytics, artificial intelligence, and night cultural tourism by connecting consumer behaviour prediction with conditional scene generation within a unified framework. Rather than treating prediction and visual generation as independent analytical tasks, the proposed architecture establishes a relationship in which predicted consumer characteristics inform the subsequent generation of tourism-related scenes. From a marketing perspective, this architecture provides a potential link between consumer intelligence and experience-oriented destination design.

The predictive results suggest that Alpha Evolution optimization can provide a useful alternative mechanism for configuring the parameters of an MLP when modelling complex consumption patterns. The proposed framework achieved an accuracy of 90.50% and an F1 score of 0.89, compared with 85.00% and 0.80, respectively, for the traditional time-series benchmark. The ablation analysis further showed a reduction in predictive performance when Alpha Evolution optimization was excluded. These results are consistent with the premise that evolutionary search can support parameter optimization in complex search spaces. Nevertheless, the findings should be interpreted as evidence under the evaluated experimental conditions rather than as evidence of universal superiority over alternative optimization strategies.

The generative results provide a complementary perspective. The conditional GAN component enables scene generation to be associated explicitly with consumer-related conditions rather than relying exclusively on unconstrained image generation. The proposed framework obtained an IS of 9.2 and an FID of 21.8, compared with lower IS values and higher FID values for the evaluated GAN benchmarks. This pattern indicates comparatively favourable computational performance in terms of image diversity and distributional similarity. However, IS and FID are technical image-generation metrics and should not be interpreted as direct measures of tourists' perceived authenticity, attractiveness, satisfaction, engagement, or experiential value. Consequently, the marketing relevance of the generated scenes requires further consumer-based validation.

The ablation findings are also relevant to understanding the functional contribution of the architecture. Removing Alpha Evolution optimization resulted primarily in a decline in predictive performance, whereas removing the conditional GAN component produced the largest reduction in the Inception Score. This pattern suggests that the two components contribute to different stages of the framework: evolutionary optimization is primarily associated with the predictive component, whereas conditional generation is particularly relevant to the diversity and quality of generated

scenes. The reduction observed when the end-to-end configuration was removed further supports the potential value of connecting prediction outputs with the subsequent generative process.

The spatial and temporal results have potential implications for tourism marketing and destination management. The concentration of consumption in shopping and cultural districts suggests that night tourism demand is spatially heterogeneous. Similarly, the temporal concentration of food and beverage consumption during the evening indicates that consumer demand varies considerably throughout the night. Such information could potentially support more context-sensitive decisions concerning service provision, promotional timing, visitor-flow management, and the design of tourism experiences. Importantly, however, the present study does not directly test whether the proposed system improves resource allocation or marketing outcomes. These applications should therefore be considered potential managerial uses rather than empirically demonstrated effects.

From a consumer behaviour perspective, the framework also offers potential opportunities for data-driven personalization. Predicted consumer characteristics could be translated into differentiated visual scenarios corresponding to distinct consumption profiles, potentially supporting more targeted communication and experience design. This approach is consistent with the increasing use of artificial intelligence to identify heterogeneous consumer preferences and develop more personalized tourism experiences. Nevertheless, the present study evaluates computational performance rather than consumers' actual responses to the generated scenarios. Future empirical research should therefore investigate whether such personalized scenes influence perceived relevance, destination image, engagement, experiential value, satisfaction, or behavioural intentions.

7. Limitations

Several limitations should be considered when interpreting the findings.

First, the empirical analysis is based on a specific consumption dataset and geographical context. Consequently, the observed performance may not generalize to cities with different tourism structures, cultural characteristics, population densities, digital infrastructures, or consumption patterns. The generalizability of the framework therefore remains to be established across broader geographical and cultural contexts.

Second, the current evaluation focuses predominantly on computational performance indicators. Accuracy and F1 score provide information about predictive performance, whereas IS and FID provide information about specific properties of generated images. However, these metrics do not directly capture the psychological or behavioural responses of tourists and consumers. In particular, high image-generation quality does not necessarily imply higher perceived authenticity, attractiveness, emotional engagement, satisfaction, or intention to visit a destination.

Third, the reported experimental results do not establish whether the observed differences between models are statistically significant. Although the proposed framework produced higher values on the reported performance measures under the evaluated conditions, stronger evidence would require repeated experiments using multiple random seeds, confidence intervals, standard deviations, and appropriate statistical comparisons. Such procedures would provide a more robust assessment of the stability of the observed differences.

Fourth, the use of SMOTE for addressing class imbalance requires careful consideration in datasets containing temporal dependencies. Although oversampling can improve the representation of minority classes, synthetic observations may modify the original distribution of the data and potentially affect temporal relationships. In the present framework, SMOTE should therefore be restricted to the training data, and its suitability should be assessed through additional sensitivity analyses.

Fifth, the framework depends on the quality, completeness, and temporal alignment of heterogeneous data sources, including remote-sensing information, points of interest, online check-ins, and historical consumption records. Missing observations, measurement errors, changes in platform usage, and differences in spatial resolution may influence model performance. These issues are particularly relevant in tourism environments, where consumer behaviour can change rapidly in response to seasonality, events, weather conditions, transportation availability, and other contextual factors.

Finally, the study does not directly examine the ethical and managerial implications of AI-enabled personalization and automated scene generation. The use of consumer-related data raises issues concerning privacy, transparency, data governance, algorithmic bias, and the potential homogenization of tourism experiences. These dimensions are particularly important when predictive systems are used to inform consumer-facing marketing and destination-management decisions.

8. Future Research

Future research should extend the proposed framework in several directions.

First, multi-city and longitudinal datasets should be incorporated to assess whether the observed performance is maintained across different tourism destinations, cultural environments, and temporal periods. Comparative studies involving cities with different levels of tourism intensity and digital infrastructure would provide stronger evidence regarding the scalability and external validity of the framework.

Second, future studies should adopt more rigorous experimental validation procedures. Multiple independent runs, different random seeds, cross-validation strategies where appropriate, confidence intervals, and statistical significance testing would allow researchers to determine whether the observed differences between the proposed framework and benchmark models are

robust. Additional comparisons with alternative evolutionary optimization algorithms and contemporary deep-learning architectures would also provide a more comprehensive evaluation of the proposed approach.

Third, future research should incorporate consumer-based evaluation of the generated scenes. Experimental studies could expose participants to AI-generated tourism scenarios and measure perceived authenticity, visual appeal, relevance, emotional response, engagement, destination image, satisfaction, and behavioural intentions. Such research would establish a direct empirical connection between computational generation quality and consumer-level outcomes.

Fourth, future research could investigate more sophisticated forms of consumer segmentation and personalization. Rather than relying exclusively on aggregate consumption characteristics, the framework could incorporate demographic, psychographic, behavioural, contextual, and experiential variables to generate more differentiated tourism scenarios. This direction would strengthen the connection between artificial intelligence, consumer insight, and personalized tourism marketing.

Fifth, future studies could integrate additional contextual variables, including weather conditions, major cultural events, transportation accessibility, social-media activity, accommodation occupancy, and real-time mobility patterns. Such variables may improve the representation of the dynamic factors that influence night-time tourism consumption and could support more context-sensitive prediction and scene generation.

Finally, future research should examine the responsible deployment of AI-generated tourism content. This includes investigating transparency mechanisms, privacy-preserving approaches, fairness and bias assessment, explainability, and human oversight. Understanding how destination managers, tourism marketers, and consumers perceive and trust AI-generated experiences will be particularly important for translating the proposed computational framework into practical tourism applications.

Taken together, the findings suggest that the integration of evolutionary optimization and conditional generative modelling represents a promising computational approach for connecting consumer behaviour prediction with AI-enabled tourism experience design. The next stage of research should move beyond technical performance and examine whether these capabilities generate measurable value for consumers, tourism organizations, and destinations.

References

1. Ali, M. M., Syed, M. H., Mohammed, Z. A., & Shahebaaz, A. (2025). "QuadAt-GAN: Quad attention enabled generative adversarial network for photorealistic image generation." *Knowledge-Based Systems*, 329, 114421.
2. Alshehri, M. S., Saidani, O., Malwi, W. A., Asiri, F., Latif, S., Khattak, A. A., & Ahmad, J. (2025). "A Hybrid Wasserstein GAN and Autoencoder Model for Robust Intrusion Detection in IoT." *Computer Modeling in Engineering & Sciences*, 143(3), 3899–3920.
3. Chen, J., Chen, W., Tang, L., Yang, X., & Li, H. (2025). "A decomposition-based evolutionary algorithm with multiple reference points strategy for multiobjective optimization." *European Journal of Operational Research*.
4. Chen, J., Li, X., Zhao, Y., Ma, Z., Han, Z., & Bao, Y. (2025). "TPCMEA: A tri-population evolutionary algorithm with adaptive stage-switching for complex CMOPs." *Applied Soft Computing*, 184, 113792.
5. Chen, W., & Zhao, H. (2025). "EMF-GAN: Efficient Multilayer Fusion GAN for text-to-image synthesis." *Computers & Graphics*, 128, 104219.
6. Guo, Z., He, J., Zhang, J., Hua, X., & Chen, Z. (2025). "Finite element model updating based on an extremely many-objective evolutionary algorithm with a modal-priority mechanism." *Mechanical Systems and Signal Processing*, 241, 113492.
7. Hu, X., Zhu, M., Xu, Z., Feng, Z., Yang, H., & Stanković, L. (2025). "Multi-task SAR image processing via GAN-based unsupervised manipulation." *Knowledge-Based Systems*, 330, 114644.
8. Jalalifar, R., Delavar, M. R., Ghaderi, S. F., & Alehashem, S. L. M. (2025). "SAGANConvLSTM: A novel spatio-temporal forecasting approach combining semivariogram-enhanced GAN and ConvLSTM for power load forecasting." *Computers and Electrical Engineering*, 128, 110718.
9. Krestenitis, M., Ioannidis, K., Vrochidis, S., & Kompatsiaris, I. (2025). "Visual to near-infrared image translation for precision agriculture operations using GANs and aerial images." *Computers and Electronics in Agriculture*, 237, 110720.
10. Li, L., Ye, P.-G., Li, Z., Yang, Z., & Zhang, Z. (2025). "Finetune and Label Reversal: Privacy-preserving unlearning strategies for GAN models in cloud computing." *Computer Standards & Interfaces*, 93, 103976.
11. Liu, C., Wang, J., You, S., Wang, D., & Yu, Z. (2025). "Machine-learning-assisted EEHEMT compact modeling of GaN HEMTs." *Microelectronics Journal*, 165, 106861.
12. Liu, W., Li, J., & Chen, S. (2025). "GSRDR-GAN: Global search result diversification ranking approach based on multi-head self-attention and GAN." *Neurocomputing*, 648, 130723.
13. Liu, Y., Hou, Y., Guo, J., & Yan, C. (2024). "The impact of environment on enhancement of tourism system resilience in China: The moderating role of digital technology." *Technological Forecasting and Social Change*, 206, 123492.
14. Lu, Y., Liu, J., & Zhang, R. (2025). "GANFR: GAN fingerprint removal network for image anti-forensics." *Knowledge-Based Systems*, 327, 114134.
15. Ma, H., Kang, X., Huang, Y., Duan, S., Li, Y., & Fang, D. (2025). "Structural transient dynamic topology optimization based on autoencoder-enhanced generative adversarial network and elitist guidance evolutionary algorithm." *Computer Methods in Applied Mechanics and Engineering*, 447, 118417.
16. Ma, X., Cui, X., Pun, M.-O., & Cheng, Z. (2026). "LTTS-

- GAN: A long-term time series generative adversarial network." *Signal Processing*, 238, 110161.
17. Nan, Y., & Wo, Y. (2025). "RIE-GAN: A retrievable image encryption method based on GAN." *Digital Signal Processing*, 163, 105202.
 18. Neri, F., Liao, L., Xue, Y., & Basgalupp, M. P. (2026). "Differentiable architecture search for GANs with rollback mechanism." *Applied Soft Computing*, 186, 114055.
 19. Nguyen, B. A., Kha, M. B., Dao, D. M., Nguyen, H. K., Nguyen, M. D., Nguyen, T. V., Rathnayake, N., Hoshino, Y., & Dang, T. L. (2025). "UFR-GAN: A lightweight multi-degradation image restoration model." *Pattern Recognition Letters*, 197, 282–287.
 20. Silva, F. B. d., Borges, L. T. M., Nepomuceno, L., & Soler, E. M. (2025). "Adaptive evolutionary algorithms with linear average convergence rate for optimization problems with Lipschitz continuous functions." *Applied Soft Computing*, 183, 113635.
 21. Song, Y., Jia, X., Jia, J., Liang, Y., & Liang, X. (2026). "3D Portrait Stylization with Adaptive Semantic Editing based on GAN Latent Codes." *Pattern Recognition*, 171, 112069.
 22. Wan, C., Fu, Y., Li, M., Lou, J., Chen, L., & Zheng, Z. (2026). "A unified gradient-flow-based GAN framework with diffusion condition guided." *Pattern Recognition*, 171, 112083.
 23. Wang, H., Ma, R., & Jiang, X. (2025). "Multiscale contextual joint feature enhancement GAN for semantic image synthesis." *Image and Vision Computing*, 161, 105637.
 24. Wang, T., Lin, T., Liu, Z., Xie, X., & Yao, S. (2026). "ADPGAN: Adaptive Differential Privacy-Preserving GAN for image privacy." *Information Fusion*, 125, 103515.
 25. Wu, Y., Gan, P., Jing, Z., & Yang, Q. (2025). "Hybrid CNN-BiLSTM-Attention Conditional GAN for power system forced oscillation localization." *International Journal of Electrical Power & Energy Systems*, 172, 111220.
 26. Xie, C., Zhang, D.-P., Su, C.-H., Chen, M.-H., & He, Y.-X. (2026). "Riding the night shift: Spatiotemporal analysis of ride-hailing's impact on urban nighttime hospitality consumption in the digital age." *International Journal of Hospitality Management*, 133, 104426.
 27. Xiong, S., Wang, H., Liao, Z., & Hashim, R. (2024). "Exploring the factors and spatial patterns of national night cultural tourism consumption agglomeration zones in China." *Heliyon*, 10(2), e24132.
 28. Yang, W., Liu, J., Liu, Y., Zhang, W., & Zheng, T. (2025). "A multiple knowledge-based evolutionary algorithm for sparse large-scale multi-objective problems." *Engineering Applications of Artificial Intelligence*, 158, 111281.
 29. Zhang, B., Luo, M., Du, Q., Yi, Z., Dong, L., Yu, Y., Feng, J., & Lin, J. (2023). "Spatial distribution and suitability evaluation of nighttime tourism in Kunming utilizing multi-source data." *Heliyon*, 9(6), e16826.
 30. Zhi, L., & Wang, W. (2025). "Research on modeling of the imbalanced fraudulent transaction detection problem based on embedding-aware conditional GAN." *Big Data Research*, 41, 100557.

Copyright: ©2026 Oleksandr P. Krupskyi, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.