Intelligent Horsetail Embroidery Production and Offshore E-commerce Economic Development Model Exploration

Rongxiang Huo

Beijing Bayi High School, Beijing, China

rongxiang.huo@bayims.cn

Keywords: Stable diffusion; Deep reinforcement learning; E-commerce platform

Abstract: As one of the less developed provinces in China, Guizhou has long faced problems such as large income disparity, poor infrastructure, and limited economic opportunities. In order to help Guizhou diversify its economy and improve the living standard of its residents, this paper proposes an intelligent horsetail embroidery production scheme based on artificial intelligence and deep reinforcement learning technology. The Stable Diffusion model is used to generate customised ponytail embroidery patterns suitable for the market demands of different countries, which reduces the design cost and improves the design efficiency. Using the improved PPO algorithm to drive the robotic arm automatic sewing technology, it significantly improves the production efficiency, reduces the manual labour input, and guarantees the product quality. Especially in the field of cross-border e-commerce, this paper analyses in detail how to use global e-commerce platforms such as Amazon, Etsy and Rakuten, combined with customised marketing strategies, to promote horsetail embroidery products to target markets such as the UK and Japan. On this basis, products that meet the aesthetics and needs of local consumers are customised according to the cultural characteristics of different countries to enhance the market competitiveness of horsetail embroidery. Meanwhile, this paper discusses the optimisation of logistics, sales and marketing strategies through cross-border e-commerce platforms to enhance the global sales and profitability of the products, to promote the in-depth integration of Guizhou's traditional handicrafts with the global market, and to realise the sustainable growth of the regional economy.

1. Introduction

Guizhou is one of China's less developed provinces, characterized by significant income disparities and lower living standards compared to more developed regions. The province faces limited access to economic opportunities and social services, further compounded by inadequate transportation infrastructure, including roads, railways, and airports, which hampers connectivity within the province and with other regions, thereby restricting trade and investment opportunities. Guizhou's economy is predominantly based on agriculture and traditional industries, with limited diversification into high-value sectors such as manufacturing and services.

A unique cultural asset of Guizhou is the ancient craft of horsehair embroidery, a technique that has been passed down for centuries. This craft is distinguished by the use of horsehair wrapped in silk threads to create embroidery, with techniques that include flat stitching, hollow stitching, raised stitching, knotting, and spiral stitching, resulting in a distinctive bas-relief effect. Horsehair embroidery pieces often convey wishes, tell stories, or encapsulate history, imbuing them with rich cultural meaning and unique artistic appeal.

The primary advocate and modernizing force behind the Shui ethnic group's horsehair embroidery is Ms. Song Shuixian, whose company markets these products through physical stores and online platforms. Her efforts have provided employment to 1,000 Shui women, 60% of whom are impoverished, thereby improving their income and quality of life. Moreover, her work has fostered the integration of local culture, tourism, and commerce.

With the rapid advancement of artificial intelligence, particularly breakthroughs in deep learning and computer vision, new approaches have emerged for designing horsehair embroidery patterns.

  DOI: 10.25236/iwmecs.2025.037

Current large models can learn from extensive paired text-image datasets to automatically generate new, creative images, significantly enhancing design efficiency, reducing costs, and catering to personalized market demands. The application of intelligent generation algorithms to horsehair embroidery design is thus of great importance. These algorithms can overcome the limitations of traditional design methods, producing more diverse and personalized designs. Additionally, the traditional manual design process is often cumbersome and time-consuming, making it difficult to meet the fast-changing demands of the market. Intelligent generation algorithms can greatly improve design efficiency, reduce costs, and be adjusted and optimized according to market demands and user preferences, meeting the needs for personalized designs.

To support Guizhou in building a more resilient and diversified economy, improving living standards, and achieving long-term growth, we propose an intelligent production scheme for horsehair embroidery. This scheme includes the use of AI to generate Shui horsehair embroidery patterns tailored for export to regions such as Japan and the United Kingdom, with customization based on the cultural preferences of these markets. Additionally, we propose the use of deep reinforcement learning-based robotic arm technology to automate the embroidery process, assisting artisans with repetitive and complex stitching tasks.

Finally, cross-border e-commerce platforms and digital marketing strategies are utilised to sell Horsetail Embroidery products globally, with targeted digital marketing campaigns utilising social media, search engine marketing and influential partnerships to attract potential customers.

2. Related Work

2.1. Text-to-image generation model

A text-to-image generation model creates images based on textual descriptions. This model enables the transformation of natural language descriptions into corresponding images, independent of existing image data, making it highly valuable for a wide range of applications.

Text-to-image generation models based on GANs (Generative Adversarial Networks) use a generator and a discriminator to learn the mapping between text and images, generating images that correspond to textual descriptions. For instance, the GAN-INT-CLS model proposed by Reed et al. includes two discriminators: one that evaluates the authenticity of the generated image and checks if it aligns with the textual description, and another that handles data interpolation [1]. Chakraborty et al. introduced the SDGAN model, which enhances high-level semantic learning using twin discriminators and contrastive loss, and integrates diverse low-level semantics in the generator through attention mechanisms and semantic conditional batch normalization [2]. However, these GAN-based models face challenges such as training instability, suboptimal image quality, and difficulty in accurately capturing and translating complex or rare textual descriptions into corresponding image features.

In contrast, text-to-image generation models based on Variational Autoencoders (VAEs) and Pre-trained Language Models (PLMs) have shown promise. DALL-E, for instance, is a model built on VQ-VAE-2 and Transformer, using VQ-VAE-2 to compress images into discrete codes, which are then interpreted and generated by a Transformer model [4]. CLIP, a multimodal pre-trained model based on contrastive learning, establishes cross-modal connections between text and images. Building on the concept of combining CLIP with other generative models, Patashnik et al. proposed the StyleCLIP model [5]. However, these VAE and PLM-based models typically have complex architectures that require significant computational resources and storage space for training and deployment [3].

Diffusion models, on the other hand, model data through a diffusion process that transitions from the data distribution to a Gaussian distribution and then denoises from the Gaussian back to the data distribution. While diffusion models offer advantages such as high sample diversity and stable training, their training and inference processes are relatively slow, which initially limited their popularity. This changed in 2019 when Nazarieh et al. introduced the Denoising Diffusion Probabilistic Model (DDPM), which not only generates high-quality images but also performs exceptionally well in tasks like image super-resolution and image restoration [6]. Okhotin et al. proposed the GLIDE model, capable of generating 1024x1024 resolution images based on any text description [7]. Google's Imagen model uses a pre-trained T5-XXL language model to encode textual information and generates images through a conditional diffusion model, producing high-fidelity and detail-rich images [8]. Rombach et al. introduced the Stable Diffusion model, which is a high-resolution, detail-rich, and semantically consistent text-to-image model [9]. Unlike traditional diffusion processes conducted in pixel space, Stable Diffusion employs the Latent Diffusion Model (LDM) paradigm, performing diffusion in a lower-dimensional latent space, significantly reducing training and inference costs [10]. Models like Stable Diffusion are noted for their ability to maintain greater diversity in generated images, achieve more realistic visuals with better detail and texture, and ensure a stable training process with less risk of mode collapse.

2.2. Deep reinforcement learning

Deep reinforcement learning (DRL) algorithms represent an emerging field in control algorithms, combining the learning capability of neural networks for low-dimensional feature representation, the powerful function approximation characteristics, and the self-learning aspect of reinforcement learning through interaction with the environment. The research on DRL began with the Deep Q-Network (DQN) proposed by Mnih et al. [11], which employs neural networks as value function approximators, enabling the learning of optimal policies in high-dimensional state and action spaces.

Xue et al. [12] proposed a Double Deep Q-Networks (DDQN)-based approach that allows mobile robots to navigate to the desired target position without colliding with any obstacles or other mobile robots. Gu [13] introduced a new Dueling Munchausen Deep Q-Network (DM-DQN) algorithm, which decouples action selection and action evaluation by decomposing the network structure into value and advantage functions, thereby accelerating convergence and enhancing generalization performance.

Although DQN-based algorithms can learn optimal value functions, they cannot directly output optimal policies. In 2015, Lillicrap et al. [14] proposed the Deep Deterministic Policy Gradient (DDPG) algorithm, which can learn optimal policies in high-dimensional state spaces for continuous control tasks. Policy gradient-based algorithms, which directly optimize the policy function, have shown superior performance compared to value function-based algorithms like DQN, attracting significant interest from researchers. The Asynchronous Advantage Actor-Critic (A3C) algorithm [15] synchronously trains the policy model (Actor) and the value model (Critic) across multiple threads, effectively utilizing computational resources and significantly improving training efficiency. Schulman et al. [16] introduced the Proximal Policy Optimization (PPO) algorithm, which combines the concepts of policy gradient and value function optimization, offering better stability and convergence performance. Other policy-based algorithms include Trust Region Policy Optimization (TRPO) [17] and Soft Actor-Critic (SAC) [18]. SAC is a reinforcement learning algorithm based on maximum entropy theory, capable of simultaneously optimizing policy and value functions, providing improved stability and exploration performance.

In summary, since the advent of DRL, there have been continuous advancements from DQN algorithms to policy gradient algorithms. Undoubtedly, DRL algorithms will become increasingly powerful in the future.

3. Relevant Technology Base

3.1. Diffusion model

The Diffusion Model is a probabilistic generative model that aims to progressively perturb data into Gaussian noise through a series of Markov chains, and then learn the reverse process to recover the original data from the noise. This process can be likened to gradually adding noise to a clear image until it becomes pure noise, and then learning how to reconstruct the clear image step by step from the noisy version.

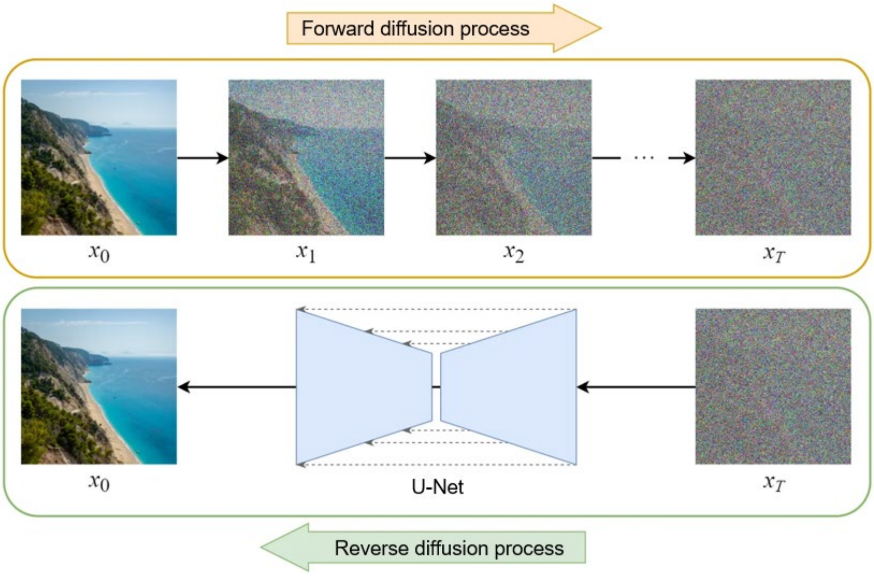

Figure 1 Schematic diagram of the Diffusion Model.

In the Diffusion Model, the forward diffusion process incrementally adds noise to a data sample x_0, generating a sequence of intermediate states $x_1, x_2, ..., x_T$ until x_T approximates isotropic Gaussian noise as shown in Fig.1. This process can be defined by a set of conditional probability distributions:

$$q(x_t \mid x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t} x_{t-1}, \beta_t \mathbf{I}) \tag{1}$$

where β_t is a small positive value representing the noise level added at each step t. Through this process, the data is gradually and completely perturbed into noise.

The goal of the reverse diffusion process is to learn how to gradually denoise the noisy data x_T to recover the original data x_0. The reverse process is typically modeled as a conditional probability distribution, where θ represents the parameters to be learned. This process can be expressed as:

$$p_\theta(x_{t-1} \mid x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \tag{2}$$

Here, μ_θ and Σ_θ represent the mean and covariance, respectively, and are typically modeled using neural networks. The training objective of the Diffusion Model is to minimize the discrepancy between the forward diffusion process and the reverse denoising process, which is achieved by minimizing the Kullback-Leibler (KL) divergence. The usual approach is to optimize the model parameters θ through backpropagation, enabling the model to better reconstruct the original data.

In practice, the training process is often optimized using the following loss function:

$$L_{\text{simple}} = \mathbb{E}_{x_0, \epsilon, t}\left[\| \epsilon - \epsilon_\theta(x_t, t) \|^2 \right] \tag{3}$$

In this context, ϵ_θ is a noise prediction model that, given a noisy sample x_t and a time step t, predicts the noise ϵ that was added.

3.2. Foundation of reinforcement learning

Reinforcement learning problems are typically modeled as either a Markov Decision Process (MDP) or a Partially Observable Markov Decision Process (POMDP), depending on whether the agent can observe all the states in the environment. These are represented by the tuples <S, A, P, R, γ> and <S, O, A, P, R, γ>, respectively, where S represents the set of states, O represents the set of observable states, A represents the set of actions, P denotes the state transition function, R denotes the reward function, and γ is the discount factor.

In reinforcement learning, the main entity is referred to as an agent. In this context, each robotic arm is considered an agent that interacts with its environment. During the reinforcement learning process, the agent observes a certain state in the environment and, based on this state, outputs an action. This action is executed in the environment, which then changes the state and provides the agent with a reward until the task is either completed or fails. The entire process from start to finish is referred to as an episode.

The types and range of actions that can be performed vary depending on the environment, and the set of valid actions is known as the action space. For instance, in the game of Go, the number of actions an agent can take is finite, whereas, in robotic arm control, the actions are often within a continuous action space.

The return is defined as the sum of all rewards from the current time step t until the end of the episode: $U_t = R_t + R_{t+1} + R_{t+2} + \ldots$ The objective of reinforcement learning is to maximize this return. Considering the uncertainty of future events, the probability of realizing a predicted reward decreases the further it is from the current time step t. In practice, a discount factor $\gamma \in [0,1)$ is applied to reduce the weight of future rewards in calculating the current return, leading to what is known as the discounted return.

$$U_t = R_t + \gamma \cdot R_{t+1} + \gamma^2 \cdot R_{t+2} + \cdots \tag{4}$$

The model consists of two main components: the reward R(s, a) provided to the agent by the environment, and the state transition probability distribution p(s'|s, a). The process of moving from the current state s to the next state s' is called a state transition, and the distribution governing this transition is known as the state transition function:

$$p_t(s' \mid s, a) = \mathbb{P}(S'_{t+1} = s' \mid S_t = s, A_t = a) \tag{5}$$

Reinforcement learning algorithms can be categorized into two major types based on whether the state transition function is known: model-based and model-free. Model-based methods involve the agent learning the state transition function to make decisions. In contrast, model-free methods neither have a defined state transition function nor attempt to estimate it; instead, they make decisions by learning either a policy network or a value function.

4. Method for Generating Horsehair Embroidery Patterns Based on Stable Diffusion

In this chapter, the pre-trained Stable Diffusion model is fine-tuned to achieve automatic generation of ponytail embroidery patterns from user-input text. By combining text input and image generation network, the model can quickly generate ponytail embroidery patterns with specific styles according to the user's needs, thus improving the efficiency and quality of ponytail embroidery pattern design.

4.1. Stable diffusion model

The structure of Stable diffusion network is shown in Fig. 2, which is mainly composed of three core components: variational self-encoder, U-Net and CLIP Text Encoder. There are three main steps from text input to image generation.

In the first step, the model receives as input a Gaussian noise matrix generated by a random function and text cues. These text cues are processed by the CLIP text encoder and transformed into text embeddings of 77x768 dimensions, where each lexical element is converted into a 768-dimensional vector. At the same time, a random latent image representation with an initial size of 64x64 is generated using seeds in the latent space.

In the second step, the denoising process is carried out using U-Net, which performs denoising operations on the random latent image representation while accepting the text embedding. U-Net progressively predicts the noise through a series of convolutional and deconvolutional steps to denoise the noisy image containing noise. In this process, the model continuously adjusts the image generation direction according to the text embedding vectors to ensure that the generated image is highly consistent with the text description.

In the third step, the VAE decoder starts working when the U-Net structure outputs the latent images. The decoder receives these latents and converts these latent representations into pixel-level image data through a series of deconvolution operations and other necessary transformations.

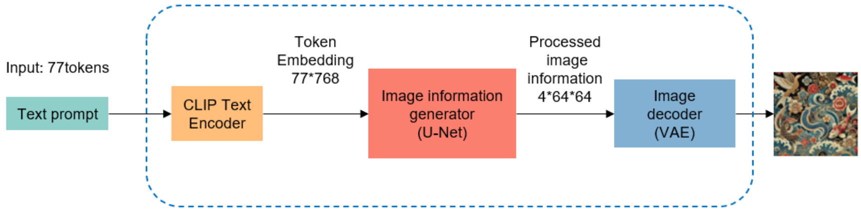

Figure 2 Stable diffusion model network structure diagram.

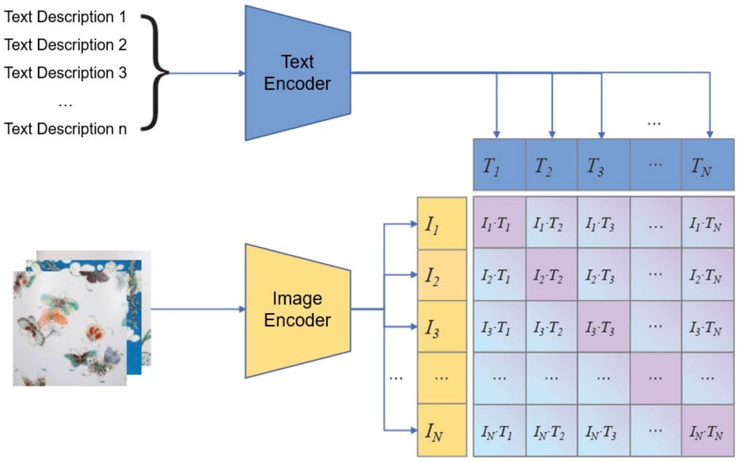

Figure 3 CLIP text encoder structure.

The CLIP text encoder architecture is shown in Fig. 3, where CLIP uses the Transformer architecture for simultaneous processing of text and image inputs. The model consists of an image encoder and a text encoder. During training, CLIP requires the model to learn to correspond text descriptions to image representations. To achieve this, CLIP uses a contrast learning approach, i.e., the model is trained by maximising the cosine similarity of matching image and text pairs. The degree of correspondence between Ia and Tb can be assessed by calculating their cosine similarity, with a higher cosine similarity indicating a closer match between text and image. Thus for N positive samples the optimisation objective is to maximise their cosine similarity and for $N^2 - N$ negative samples the optimisation objective is to minimise their cosine similarity. Once training is complete, the CLIP model can optionally be fine-tuned for a specific dataset task or downstream task. Typically fine-tuning uses supervised learning methods on a specific task to adjust model parameters.

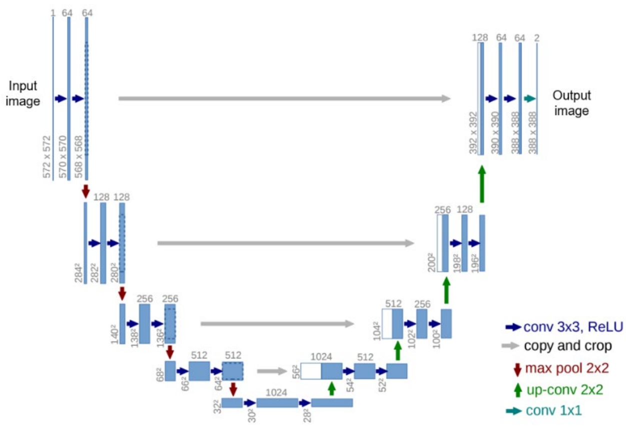

Figure 4 U-Net structure.

The U-Net network structure, shown in Fig. 4, is a symmetric network structure consisting of a left compression path (encoder) and the right expansion path (decoder) together. The encoder consists of a convolutional layer and a maximum pooling layer that convolves the input image multiple times to extract key features and down-samples the feature map by pooling the layer to reduce its dimensionality . The decoder uses up-sampling and inverse convolution to recover the information lost by the encoder. The encoder and decoder exchange shallow and deep feature information through hopping connections and increase the nonlinearity with a ReLU activation function after each convolution. In the four-layer architecture of U-Net, the first layer receives a 572x572x1 image and obtains a feature map of 64x568x568 by two 3x3 convolutions, followed by 2x2 pooling to downscale to 284x284x64, and then each layer continues to extract and downscale the image features by two convolutions and one down-sampling.

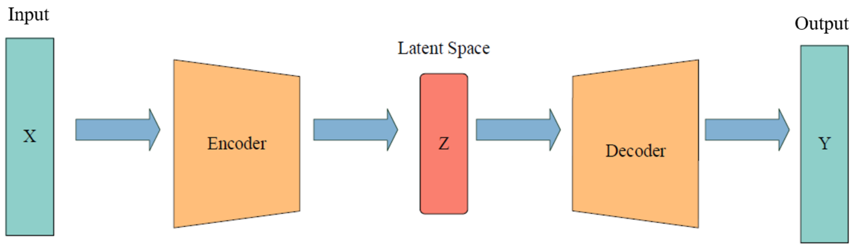

Figure 5 VAE structure.

The basic structure of the Variational Auto-Encoder (VAE) is shown in Fig. 5. The VAE consists of an encoder and a decoder connected through an intermediate representation called the latent space. The encoder converts the input data into a simplified latent representation through a series of convolutional layers, gradually reducing the spatial size of the image and increasing the feature depth. The latent space receives the output of the encoder, i.e., the statistical parameters (mean and variance) of the data, and provides it to the decoder. The decoder task is to reconstruct an image similar to the original data from the representation in the latent space using an inverse convolution or up-sampling layer to recover its original dimensions.

4.2. Horsetail embroidery pattern generation results

Based on the above Stable Diffusion's text-to-image generation model, we test the generation of horsetail embroidery patterns using cultural features of different countries as semantic inputs. Fig. 6(a) illustrates ponytail embroidery with Japanese characteristics, including intricate floral and wave motifs, cherry blossoms, koi and cranes, combined with rich and vibrant colours. Fig. 6(b) shows horsetail embroidery with British characteristics, including roses, lilies, shields and lions, with Victorian motifs adding a historical touch, and rich tones of crimson, blue, gold and green reflecting the harmonious fusion of Chinese and British artistic traditions.

(a) (b)

Figure 6 (a) Horsetail embroidery with Japanese characteristics. (b) Horsetail embroidery with British characteristics.

5. Automatic Sewing with Robotic Arm Based on Improved PPO

In recent years, humanoid robot technology has developed rapidly, of which the operation of robotic arms is a key part of its application. With high precision and flexibility, robotic arms are widely used in the field of fine manipulation. In traditional crafts such as horsetail embroidery, the sewing process is cumbersome and time-consuming, and the craftsmen are easily fatigued, while deep reinforcement learning-based robotic arm automatic sewing technology effectively solves this problem. Deep Reinforcement Learning enables the robotic arm to autonomously learn complex sewing tasks through environmental interactions and strategy optimisation, precisely executing stitches and adapting to different patterns and materials. This technology not only improves efficiency, but also reduces the workload of the embroiderer while maintaining high quality embroidery results.

5.1. Proximal Policy Optimization

Proximal Policy Optimization (PPO) is a commonly used policy optimization algorithm in reinforcement learning, which improves the stability and efficiency of the algorithm by limiting the magnitude of policy updates. The core idea of PPO is to update the policy parameters by optimising the objective function, so that the policy can achieve a better performance while keeping small changes.

Traditional policy gradient methods (e.g., REINFORCE or Actor-Critic), although they can directly optimise the policy, they may face the following problems: High variance in the policy update may lead to an unstable learning process. Besides, if the policy update is too large, it may cause the policy to deviate from a known good policy, leading to performance degradation.

As illustrated in Figure 7, the PPO algorithm follows an Actor-Critic architecture centered on a trajectory buffer. The process begins with the agent collecting transition tuples (s_t, a_t, r_t) from the environment, which are then used to concurrently update the two networks. The Actor is optimized via a surrogate loss $\mathcal{L}_{surrogate}$ that incorporates the probability ratio between the new and old policies, weighted by the Generalized Advantage Estimate (GAE) $\hat{A}_t$ to ensure stable policy shifts. Simultaneously, the Critic is refined by minimizing the value loss $\mathcal{L}_{value}$, aligning the predicted state-value $V_\eta(s_t)$ with the observed returns. This iterative feedback loop enables the model to balance exploration and exploitation while maintaining training stability.

The main purpose of PPO is to control the magnitude of policy changes during policy updates to ensure the stability of the updates. It does this by introducing a clipping mechanism to limit the magnitude of each policy update to avoid excessive policy changes. Specifically, the goal of PPO is to maximise the following objective function:

$$L^{\text{CLIP}}(\theta) = \mathbb{E}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right] \tag{6}$$

where $r_t(\theta)$ is the strategy ratio, defined as the ratio of the new strategy to the old strategy in the same state and action:

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)} \tag{7}$$

A^t is the Advantage Function, which measures how much better than average it is to choose a particular action in the current state. ϵ is a hyperparameter that controls the magnitude of the strategy update. The core of the clipping mechanism is to avoid excessive changes in strategy updates by limiting the strategy ratio $r_t(\theta)$ to [1-ϵ,1+ϵ]. This approach ensures that the strategy does not deviate too far during the update, thus maintaining the stability of the update process.

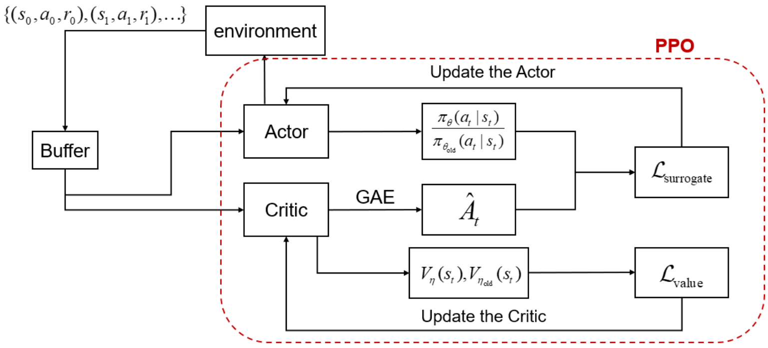

Figure 7 PPO algorithm block diagram.

5.2. Improved PPO

Considering the performance of the PPO algorithm in the simulation experiments of the robotic arm embroidery task, we try to integrate the idea of maximum entropy in the PPO algorithm to ensure sufficient exploration of the intelligences so that the algorithm is fast and stable. For the convenience of elaboration, the improved algorithm is called PPO_IMPROVED algorithm in this paper.

The policy function of the PPO_IMPROVED algorithm can be expressed as:

$$L_t^{IMPROVED}(\theta)=\hat{E}_t\left[L_t^{CLIP}(\theta)-c_1L_t^{VF}(\theta)+c_2S[\pi_\theta](s_t)\right] \tag{8}$$

c_1,c_2 are the coefficients, S denotes the value of entropy, and $L_t^{VF}(\theta)$ is an estimate of the value function. The value function shear is used in the PPO_IMPROVED algorithm with the following expression:

$$L_t^{VF}(\theta)=\max\left[\left(V_\theta(s_t)-V_t^{targ}\right)^2,\left(clip\left(V_\theta(s_t),V_\theta(s_{t-1})-\sigma,V_\theta(s_{t+1})+\sigma\right)-V_{targ}\right)^2\right] \tag{9}$$

In addition to this, the PPO_IMPROVED algorithm uses a decayable learning rate:

$$\gamma_n=\gamma(1-step_t/step_a) \tag{10}$$

γ_n denotes the current learning rate, $step_t$ is the number of current training steps, and $step_a$ is the total number of training steps. The way of learning rate decaying can improve the stability of the algorithm in the later stage of training to a certain extent.

5.3. Experiment results

The PPO_IMPROVED algorithm is applied to the same simulation environment as the DDPG, TD3, PPO and SAC algorithms to simulate the embroidery of a six-degree-of-freedom robotic arm. Figure 8 shows the comparison of the results of the five algorithms, DDPG, TD3, PPO, SAC and PPO_IMPROVED, in the embroidery simulation of the robotic arm. From the figure, we can find that the DDPG algorithm is extremely unstable during the training process, and the gain curve of the algorithm has strong fluctuations. td3 algorithm has improved the stability of the DDPG algorithm, but there are still some fluctuations in the training curve. the SAC algorithm has the first convergence speed in the whole training process, and the stability of the algorithm is not bad, but the cumulative rewards of the intelligences are relatively low. ppo algorithm has the best performance in control stability and rewards, and it has the best performance in control stability and reward. The PPO algorithm is very good in terms of control stability and rewards, but the convergence speed is not very fast. The PPO_IMPROVED algorithm has partially improved on this, increasing the convergence speed of the algorithm as well as the gain of the intelligences, and is the best of the five algorithms.

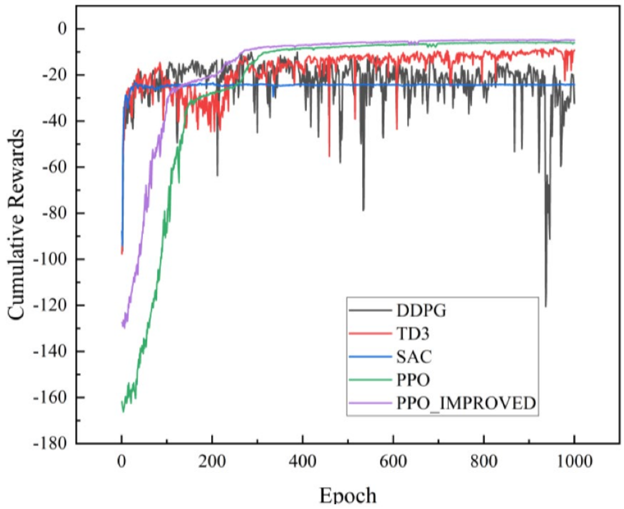

Figure 8 Cumulative returns of five algorithms.

Finally, by deploying the above deep reinforcement learning algorithms into the robotic arm, the integrated AI-controlled robotic arm can help the embroiderer complete repetitive and complex sewing tasks. The AI-based robotic arm control system can characterise the digital embroidery pattern and control the movement of the robotic arm, and the embedded AI needle can ensure accurate and consistent sewing results with the embroiderer. Reducing the physical burden on the embroiderer improves efficiency. Fig. 9 shows the robotic arm sitting on top of a piece of fabric, equipped with precision sensors and actuators with embedded sensors and AI algorithms to accurately guide the needle's movement.

Figure 9 The robotic arm sits above a piece of fabric.

6. Economic development model of Horsetail embroidery based on cross-border e-commerce platform

6.1. Cross-border e-commerce platform selection and analysis of the UK and Japanese markets

The choice of cross-border e-commerce platforms is key to driving the globalisation of horsetail embroidery sales, especially in the UK and Japan, two countries with different cultural backgrounds and market demands. As a European economic powerhouse, the UK has a mature e-commerce market and high purchasing power, while Japan is one of the developed countries in Asia, focusing on culture and fine consumer goods. For these two markets, a reasonable choice of platforms can help to promote horsetail embroidery to the international stage, expand sales channels and enhance its cultural value.

Amazon, as a leading global platform, has a large market share in both the UK and Japan. Its logistics system and large user base can help horsetail embroidery products quickly enter consumers' homes. The UK market has a high demand for home décor and handicrafts, and the cultural connotation of horsetail embroidery fits this demand. In the Japanese market, Amazon's reach makes it ideal for promoting exquisite handicrafts, and through its rating mechanism and recommendation algorithm, the characteristics of horsetail embroidery can be spread more widely, while Etsy is suitable for attracting users who are interested in traditional culture and customised products. In the UK and Japan, Etsy provides an ideal platform for horsetail embroidery, incorporating cultural storytelling to promote and enhance the added value of the product. In addition, Etsy supports personalised promotions through social media and advertising campaigns to expand brand awareness, while Rakuten, a local e-commerce platform in Japan, helps horsetail embroidery products respond to the needs of local consumers thanks to its localised services. By incorporating Japanese elements such as cherry blossoms and koi carp into the pattern design, horsetail embroidery can better attract the love of Japanese consumers. All in all, by utilising platforms such as Amazon, Etsy and Rakuten, horsetail embroidery is able to achieve multi-channel sales in the UK and Japan, enhance its brand influence, and achieve long-term market expansion and brand accumulation through cultural promotion.

6.2. Product customisation strategies for the UK and Japanese markets

In promoting the entry of horsetail embroidery products into the British and Japanese markets, the development of customisation strategies is particularly important, which involves not only an in-depth understanding of the cultural differences between the two countries, but also an accurate grasp of consumers' aesthetic preferences and market demand. The UK and Japan have very different cultural backgrounds, but both markets show a high degree of recognition of traditional crafts, so in customised design, local cultural symbols and traditional values should be fully integrated to enhance the attractiveness and market competitiveness of horsetail embroidery products.

In the UK market, consumers tend to favour goods with a sense of history and artistic value, so horsetail embroidery products can be integrated with classic British cultural elements, such as roses, lilies, shields, lions and other symbolic motifs, which not only carry the history and culture of the UK, but also can be cleverly combined with the craftsmanship characteristics of horsetail embroidery. At the same time, British consumers prefer calm and elegant colours, especially deep red, deep blue, gold and other retro shades, which can effectively enhance the artistic sense and high-end positioning of horsetail embroidery products. In addition, the UK market has a high demand for high-end home decorations, so the design can combine the traditional craftsmanship of horsetail embroidery with modern home styles, and launch unique embroidered wall hangings, cushions and other home furnishings to meet consumer demand for unique decorations.

Compared to the UK market, Japanese consumers are more focused on simplicity and natural beauty, which coincides with the delicate craftsmanship displayed in horsetail embroidery. The reverence and respect for nature in Japanese culture has led to natural elements such as cherry blossoms, koi and waves being favoured in product design. Therefore, when customising horsetail embroidered products for the Japanese market, these elements, which symbolise the beauty of nature, can be combined with the embroidery techniques of the aquatic culture to create products that are both aesthetically pleasing to the Japanese and rich in the characteristics of Chinese minority culture. In addition, Japanese consumers have strict requirements for the practicality and detail design of products, so product design should focus on detail processing and a sense of high quality, such as through the soft pink, blue and other light colours to enhance the visual appeal of the product. At the same time, Japanese consumers respect the spirit of craftsmanship, so in the packaging and publicity of the product, special emphasis can be placed on the handmade skills and heritage value of horsetail embroidery to meet the preference of Japanese consumers for high-quality handicrafts.

6.3. Profitability and economic modelling

When constructing a cross-border e-commerce profitability model for horsetail embroidery products, the rational design of the cost model and revenue projections are crucial. The core of the profit model is to ensure that the products are competitive in the international market while achieving long-term economic sustainability. Therefore, accurately measuring the pricing strategy, cost sharing, and profitability of the product is a prerequisite to ensure the success of this business model.

Firstly, as a high value-added handicraft, product pricing of horsetail embroidery should take into account a number of factors such as production costs, logistics costs, platform service fees and marketing costs. According to these cost elements, the sales pricing should ensure profit maximisation under the premise of ensuring consumer acceptance. Assuming that the selling price of the product is P, and its cost components mainly include material cost C_m, labour cost C_a, platform service fee F_p and logistics fee L_f, the total cost of each product C_t can be expressed as follows:

$$C_t = C_m + C_a + F_p + L_f \tag{11}$$

The gross profit G is then the difference between the selling price and the total cost, i.e.:

$$G = P - C_t \tag{12}$$

This formula suggests that profits depend directly on the difference between pricing P and costs C_t. The price sensitivity of the cross-border e-commerce market must be taken into account when developing a pricing strategy, especially in highly competitive international markets, where too high a pricing may dampen demand, while too low a pricing may lead to a narrowing of profit margins. Therefore, ideal pricing should ensure both product competitiveness and profit margins.

Second, in cross-border e-commerce, logistics costs L_f have a significant impact on profitability. Cross-border transport costs are high and are affected by the international economic environment, tariff policies and logistics conditions in the destination country. Therefore, in order to optimise profitability, enterprises need to establish long-term relationships with logistics providers and reduce logistics costs through bulk shipments and efficient delivery solutions. In addition, platform service fee F_p is also an influencing factor. The commission structure of different e-commerce platforms varies, usually accounting for 10-15 per cent of the selling price of the product, so it is necessary to fully weigh the cost of platform selection against the benefits of platform traffic.

While ensuring profitability, break-even analysis is a key step in developing a business strategy. The break-even point is calculated based on the fixed costs F and the profit per unit of product G, i.e.:

$$Q_{\text{break-even}} = \frac{F}{G} \tag{13}$$

where $Q_{\text{break-even}}$ denotes the minimum volume of sales required to reach break-even. If the firm's monthly fixed costs (e.g., employee wages, equipment maintenance, and platform rental) are F = \$15,000, and assuming a profit of G = \$45 per product, it would need to sell approximately $Q_{\text{break-even}} = \frac{15,000}{45} = 334$ products to reach the break-even point.

Further, the sustainability analysis based on this profitability model should also be used to project future revenues in conjunction with market sales expectations and demand growth. Considering the gradual expansion of the market, it is expected that the annual sales volume will increase every year as the brand awareness increases, and assuming an annual growth rate of r, the annual revenue, R, can be calculated based on the compound growth formula:

$$R_t = R_0 (1 + r)^t$$

where R_0 is the initial annual revenue, t is the year, and r is the annual growth rate. This formula allows revenue growth to be projected for the next five years and long-term profitability to be assessed.

In revenue growth forecasting, cost control is key to achieving profitability sustainability. By optimising production processes, reducing logistics and platform costs, and increasing the level of automation in the sewing process, unit costs can be effectively reduced, thereby increasing gross margins. Assuming that the company is able to reduce its costs by c% per annum through technological advancement in the course of revenue growth, the growth in net profit per annum can be expressed by the following formula:

$$P_t = (R_t - C_t) = R_0 \times (1+r)^t - C_0 \times (1-c)^t \tag{14}$$

Where, C_t is the total cost in year t, C_0 is the initial cost, and P_t is the net profit in year t. By reasonably controlling costs and driving revenue growth, firms can break even and increase profitability in a relatively short period of time. Sustainability analysis should also consider the saturation of the market and the diversification of consumer demand. In future market expansion, horsetail embroidery products need to rely not only on the core markets such as the UK and Japan, but also consider further development of other potential markets. Market expansion will help diversify business risks and ensure profitability stability in different economic environments. Combined with long-term market dynamics and brand expansion strategies, the Company can effectively enhance its ability to withstand market volatility and thus achieve sustainable revenue growth.

7. Conclusion

This study proposes a model integrating intelligent production and cross-border e-commerce to boost Guizhou's economy via its horsetail embroidery heritage. In intelligent production, the fine-tuned Stable Diffusion model—combined with CLIP Text Encoder—generates market-tailored patterns (e.g., Japanese-style with cherry blossoms, British-style with roses), cutting design costs and meeting diverse demands. The improved PPO (PPO_IMPROVED) algorithm drives robotic arm sewing, outperforming DDPG, TD3, SAC, and original PPO in convergence, stability, and rewards, reducing manual work and ensuring quality.

For market expansion, targeted strategies for the UK and Japan include platform selection (Amazon for coverage, Etsy for cultural customization, Rakuten for Japanese localization) and product customization matching local aesthetics. The profitability model clarifies cost-revenue relationships; via cost optimization (e.g., bulk logistics) and sales growth, short-term break-even and long-term sustainable revenue are achievable.

Overall, this model integrates Guizhou's traditional crafts with AI and global markets, offering a path for cultural heritage inheritance, economic diversification, and regional sustainable development. Future research could expand to more markets and deepen AI-industrial chain integration.

References

[1] Ye S. Text Generation Image about Gan[C]//2021 International conference on Smart Technologies and Systems for Internet of Things (STS-IOT 2021).Atlantis Press, 2022:119-123.

[2] Chakraborty T, KS U R, Naik S M, et al. Ten years of generative adversarial nets(GANs): a survey of the state-of-the-art[J]. Machine Learning: Science and Technology, 2024, 5(1): 011001.

[3] Zhao W X, Liu J, Ren R, et al. Dense text retrieval based on pretrained language models: A survey[J]. ACM Transactions on Information Systems, 2024, 42(4): 1-60.

[4] Alfano M, Abedin E, Reimann R, et al. Now you see me, now you don't: an exploration of religious exnomination in DALL-E[J]. Ethics and Information Technology, 2024, 26(2): 1-13.

[5] Zhang B, Zhang P, Dong X, et al. Long-CLIP: Unlocking the Long-Text Capability of CLIP[J]. arxiv preprint arxiv:2403.15378, 2024.

[6] Nazarieh F, Feng Z, Awais M, et al. A Survey of Cross-Modal Visual Content Generation[J]. IEEE Transactions on Circuits and Systems for Video Technology, 2024.

[7] Okhotin A, Molchanov D, Vladimir A, et al. Star-shaped denoising diffusion probabilistic models[J]. Advances in Neural Information Processing Systems, 2024, 36.

[8] Hu H, Chan K C K, Su Y C, et al. Instruct-Imagen: Image generation with multi-modal instruction[J]. arxiv preprint arxiv:2401.01952, 2024.

[9] Li Y, Wang H, et al. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds[J]. Advances in Neural Information Processing Systems, 2024, 36.

[10] Zhao P, Xu P, Qin P, et al. LAKE-RED: Camouflaged Images Generation by Latent Background Knowledge Retrieval-Augmented Diffusion[J]. arxiv preprint, arxiv:2404.00292, 2024.

[11] Mnih V, Kavukcuoglu K, Silver D, et al. Human-level control through deep reinforcement learning[J]. Nature,2015,518(7540):529.

[12] Xue X, Li Z, Zhang D, et al. A deep reinforcement learning method for mobile robot collision avoidance based on double dqn[C]//2019 IEEE 28th International Symposium on Industrial Electronics (ISIE). IEEE, 2019: 2131-2136.

[13] Gu Y, Zhu Z, Lv J, et al. DM-DQN: Dueling Munchausen deep Q network for robot path planning[J]. Complex & Intelligent Systems, 2022: 1-14.

[14] Lillicrap T P, Hunt J J, Pritzel A, et al. Continuous control with deep reinforcement learning[J]. US Patent, 2020, 15(217,758).

[15] Leng J, Fan S, Tang J, et al. M-A3C: A Mean-Asynchronous Advantage Actor-Critic Reinforcement Learning Method for Real-Time Gait Planning of Biped Robot[J]. IEEE Access, 2022, 10: 76523-76536.

[16] Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms[J]. arXiv preprint arXiv:1707.06347, 2017.

[17] Schulman J, Levine S, Abbeel P, et al. Trust region policy optimization[C] //International conference on machine learning. PMLR, 2015: 1889-1897.

[18] HAARNOJA, T., ZHOU, A., ABBEEL, P., et al. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor[C]//International conference on machine learning. PMLR, 2018: 1861-1870.