

ERL-MPP: Evolutionary Reinforcement Learning with Multi-head Puzzle Perception for Solving Large-scale Jigsaw Puzzles of Eroded Gaps

Xingke Song¹, Xiaoying Yang¹, Chenglin Yao¹, Jianfeng Ren^{1,2*}, Ruibin Bai^{1,2}, Xin Chen³, Xudong Jiang⁴

¹School of Computer Science, University of Nottingham Ningbo China, China
²Nottingham Ningbo China Beacons of Excellence Research and Innovation Institute, University of Nottingham Ningbo China, China

³Intelligent Modelling & Analysis Group, School of Computer Science, University of Nottingham, Nottingham, UK

⁴School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore

{Xingke.Song, Xiaoying.Yang, Chenglin.Yao, Jianfeng.Ren, Ruibin.Bai}@nottingham.edu.cn, Xin.Chen@nottingham.ac.uk, EXDJiang@ntu.edu.sg

Abstract

Solving jigsaw puzzles has been extensively studied. While most existing models focus on solving either small-scale puzzles or puzzles with no gap between fragments, solving large-scale puzzles with gaps presents distinctive challenges in both image understanding and combinatorial optimization. To tackle these challenges, we propose a framework of Evolutionary Reinforcement Learning with Multi-head Puzzle Perception (ERL-MPP) to derive a better set of swapping actions for solving the puzzles. Specifically, to tackle the challenges of perceiving the puzzle with gaps, a Multi-head Puzzle Perception Network (MPPN) with a shared encoder is designed, where multiple puzzlet heads comprehensively perceive the local assembly status, and a discriminator head provides a global assessment of the puzzle. To explore the large swapping action space efficiently, an Evolutionary Reinforcement Learning (EvoRL) agent is designed, where an actor recommends a set of suitable swapping actions from a large action space based on the perceived puzzle status, a critic updates the actor using the estimated rewards and the puzzle status, and an evaluator coupled with evolutionary strategies evolves the actions aligning with the historical assembly experience. The proposed ERL-MPP is comprehensively evaluated on the JPLEG-5 dataset with large gaps and the MIT dataset with large-scale puzzles. It significantly outperforms all state-of-the-art models on both datasets.

Introduction

Jigsaw puzzles are popular entertainment and intellectual challenges. Solving jigsaw puzzles has been related to a wide range of applications, *e.g.*, geometric analysis (Harel, Shahar, and Ben-Shahar 2024), boundary information understanding (Gallagher 2012), and image semantic understanding (Song et al. 2023a). In particular, automatic puzzle solving has been extensively studied in archaeology (Nada and Jan 2015) and cultural heritage restoration (Derech, Tal, and Shimshoni 2021), which greatly releases the burden of archaeologists. The developed techniques have been applied beyond puzzle reassembly, *e.g.*, self-supervised learn-

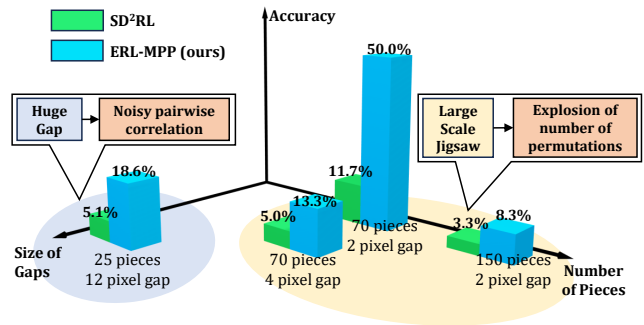


Figure 1: Huge gaps and a large amount of fragments impose great challenges on puzzle solvers. The proposed ERL-MPP significantly outperforms the previous best performing SD²RL (Song et al. 2023a) on four different types of puzzles, *e.g.*, puzzles with as large as 12-pixel gaps, or puzzles with as many as 150 pieces of fragments.

ing (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016; Ren et al. 2023), video spatial understanding (Yang et al. 2022), fine-grained visual classification (Du et al. 2020), and image super-resolution (Ma et al. 2021).

Most methods focus on solving puzzles with no gaps or very small gaps between fragments (Son, Hays, and Cooper 2019; Bridger, Danon, and Tal 2020; Li et al. 2022), and utilize the boundary information to infer the pairwise relations between fragments. However, in real-world applications, there are often gaps on artifacts due to damage or erosion (Paumard, Picard, and Tabia 2020; Song et al. 2023a), making the boundary similarity unavailable. This paper tackles the problem of solving Large-scale Jigsaw Puzzles of Eroded Gaps (L-JPEG). Challenges are two-fold. 1) Visual perception of fragments under gaps. Large gaps make the methods based on the boundary information ineffective, resulting in noisy estimation of relations between adjacent fragments. 2) A huge search space for optimization strategy. The number of permutations increases exponentially with the number of puzzle pieces, coupled with the uncertainty in pairwise relations, making the puzzle re-

*Corresponding Authors.

assembly computationally demanding.

In this paper, we propose a framework of Evolutionary Reinforcement Learning with Multi-head Puzzle Perception (ERL-MPP) to derive a better set of swapping actions to assemble the puzzle, leading to the significant performance gain as shown in Fig. 1. Firstly, to perceive visual clues from puzzles effectively, a Multi-head Puzzle Perception Network (MPPN) is designed as the perception module, where three puzzlet perception heads excavate the local features at multiple scales by comparing fragments in local neighborhoods, and a discriminator head globally evaluates whether the puzzle has been correctly reassembled. The feature encoder, decoder, and discriminator form a Generative Adversarial Network (GAN), where the discriminator evaluates the quality of decoded image as a reassembled puzzle and helps refine the encoder to capture the critical features in assessing the puzzle status. The puzzlet perception heads and the discriminator head share the same feature encoder to not only reduce the network complexity but also produce robust features.

Secondly, to tackle the challenges of assembling a large number of puzzle pieces, an Evolutionary Reinforcement Learning (EvoRL) agent is designed to select a set of fragment-swapping actions based on both the current visual perception and the past reassembly experience stored in an agent. RL agents have demonstrated good scalability (Arulkumar et al. 2017), and evolutionary strategies can reduce redundant exploration, thereby improving sample efficiency and diversity (Zhang et al. 2024b). The proposed EvoRL consists of an actor, a critic, and an evaluator. While the actor predicts the probability of feasible actions from a given action space, the critic evaluates the current puzzle status and updates the actor based on the visual perception from the MPPN and the reward collected from the environment. The evaluator assesses the quality of actions evolved using techniques such as crossover with the action population from past reassembly experience and natural mutations, which helps discover high-performance actions, facilitating the agent to fast converge to an optimal solution. Besides the actions of swapping two fragments (Song et al. 2023a), the agent explores the actions of swapping three fragments and swapping two 2×2 puzzlets. During training, the agent initially exploits larger size swapping actions to explore the action space more efficiently with the help of the evolutionary strategies to avoid a local optimum, and gradually utilizes smaller size swapping actions to focus on assembling the puzzle perfectly. A set of rewards are designed to provide clear training guidance for the agent, promoting the correct fragment placement and perfect puzzle reassembly. By utilizing the multiple perceptions of the MPPN and the effective exploration of the EvoRL, the proposed model better explores and exploits the large action space to derive a better set of swapping actions to perfectly reassemble the puzzle.

Our contributions can be summarized as follows. 1) The proposed MPPN perceives both locally adjacent fragments in small neighborhood through the puzzlet perception heads and the global puzzle status through the discriminator head. 2) The proposed EvoRL parameterizes the action probabilities from the action space and utilizes the evolutionary strategies to evolve series of best swapping trajectories, alle-

viating the problem of enumerating and evaluating all possible swapping actions in the value-based RL agents (Song et al. 2023a), which effectively explores the large action space for large-scale puzzles. 3) The proposed ERL-MPP consistently and significantly outperforms the state-of-the-art models on the JPLEG-5 dataset (Song et al. 2023a) with large gaps and the MIT dataset (Bridger, Danon, and Tal 2020) with large-scale puzzles.

Related Work

Puzzle solving often involves visual perception of puzzles and puzzle reassembly based on the perceived information.

Visual Perception of Puzzles. Traditional methods mainly handle puzzles with no gap or very small gaps between fragments, and utilize the boundary information to reliably estimate the pairwise relations between fragments, *e.g.*, Handcrafted features such as boundary shapes (Zhang et al. 2015), color contrast (Son, Hays, and Cooper 2019; Yan et al. 2021), and boundary similarities (Sholomon, David, and Netanyahu 2013; Paikin and Tal 2015) are often utilized to describe adjacent fragment pairs. Deep neural networks have also been developed to extract boundary features (Khoroshiltseva et al. 2022; Li et al. 2022).

Recent studies focus on solving puzzles with gaps, where boundary information is not available due to relatively large gaps, *e.g.*, 2-4 pixels (Bridger, Danon, and Tal 2020), 12 pixels (Song et al. 2023b,a) and 48 pixels (Paumard, Picard, and Tabia 2020; Song et al. 2023a,b). In this case, image semantics is often utilized to exploit pairwise relations (Paumard, Picard, and Tabia 2020; Song et al. 2023a,b), serving as important visual clues as in other applications (Li et al. 2024; Zhang et al. 2024a, 2023). The gaps between adjacent pieces are often inpainted by using GAN (Bridger, Danon, and Tal 2020; Khoroshiltseva et al. 2022). But it is difficult to inpaint the large gaps of 12 pixels or 48 pixels. Deep neural networks are often utilized to estimate the adjacency relations (Paumard, Picard, and Tabia 2018, 2020; Song et al. 2023a,b). To predict the correlation between the central fragment and adjacent ones, Paumard, Picard, and Tabia (2020) developed a Siamese Neural Network. In SD²RL, Song et al. (2023a) designed a set of Siamese Discriminant Networks to infer the pairwise relations. Besides pairwise relations, puzzlets formed by more adjacent pieces are evaluated by Puzzlet Discriminant Network (Song et al. 2023b). These existing methods often focus on local pairwise features, while global puzzle perception could also help. This paper presents a multi-head puzzle perception network to perceive both local and global puzzle placement status.

Puzzle Reassembly. Puzzle reassembly strategies can be broadly divided into two groups. 1) Methods for solving puzzles with no gaps or small gaps (Son, Hays, and Cooper 2019; Bridger, Danon, and Tal 2020), where adjacency between fragments can be reliably estimated, and maximizing adjacency generally leads to correct placements (Cho, Avidan, and Freeman 2010; Gallagher 2012). These methods focus on accelerating the solution of large-scale puzzles, *e.g.*, greedy tree-based method (Bridger, Danon, and Tal 2020), genetic algorithm (Sholomon, David, and Netanyahu 2013), and loop constraints (Son, Hays, and Cooper

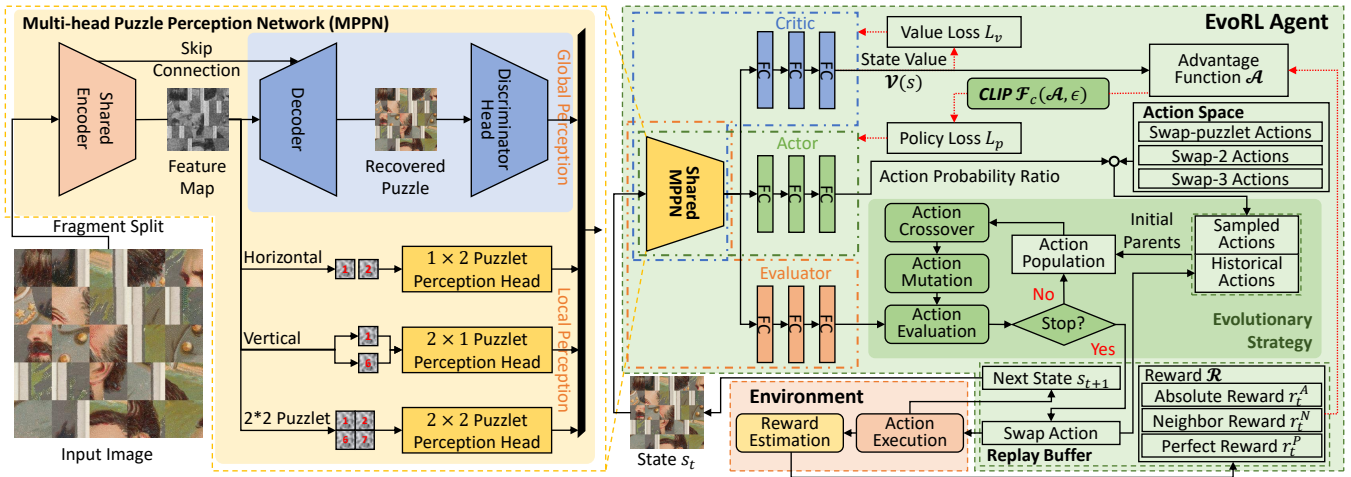


Figure 2: Block diagram of the proposed ERL-MPP for solving L-JPEG problems. The shared MPPN globally perceives the puzzle through a discriminator head and locally perceives it through three puzzlet perception heads. The EvoRL agent determines an optimal sequence of swapping actions till perfectly reassembling the puzzle, where an actor recommends swapping actions from a large action space of Swap-2, Swap-3, and Swap-Puzzlet actions, a critic estimates the state value and updates the actor based on the visual perception from the MPPN and the estimated reward, and an evaluator assesses and selects the most suitable action after evolutionary operations such as crossover and mutation. A set of rewards considering fragment placement, pairwise adjacency and perfect reassembly are designed to guide the training of the agent.

2014; Son, Hays, and Cooper 2019). 2) Methods for solving puzzles with large gaps (Paumard, Picard, and Tabia 2020; Song et al. 2023a,b), where the adjacency between fragments can’t be reliably estimated due to the gaps. Paumard, Picard, and Tabia (2020) formulated puzzle reassembly as a shortest-path optimization problem, and solved it by Dijkstra’s algorithm with branch-cut. PDN-GA (Song et al. 2023b) utilizes a genetic algorithm to derive the reassembly sequence. Song et al. (2023a) developed a Deep Q-network for puzzle solving, but it is difficult to evaluate all the feasible actions, especially when the puzzle has many fragments.

Proposed ERL-MPP

Overview of Proposed ERL-MPP

The proposed ERL-MPP tackles the L-JPEG problems, where both the gaps and the large number of puzzle pieces pose great challenges. The block diagram is shown in Fig. 2. The proposed shared MPPN consists of a GAN-like network head to perceive the global puzzle status and three puzzlet perception heads to perceive local puzzle status. The discriminator head and the puzzlet perception heads share the same encoder to generate a comprehensive representation through joint feature learning. The encoder-decoder architecture together with the discriminator is designed to derive robust features focusing on the image semantics.

The proposed EvoRL agent includes three major components. 1) An actor to estimate the probability ratio of swapping actions in a given action space, and recommend candidate actions for the evolutionary algorithm to form high quality swap-based search trajectories with the assistance of historical experience. 2) A critic to estimate the state value based on the visual perception from the MPPN and the re-

ward evaluated from the puzzle placement, and update the critic through the critic loss and the actor through the policy loss. 3) An evaluator to assess the swapping actions in a rollout fashion after crossover and mutation, and to select the most suitable one. It shares the same architecture as the critic and is soft-updated by the critic (Wei et al. 2021), learning from the critic while maintaining independence to prevent potential overfitting. The evolution helps explore a wider range of combinations of swapping actions, potentially discovering high-performing actions that the actor might not find on its own. Compared with value-based reinforcement learning (Song et al. 2023a), the EvoRL better balances the exploration and exploitation in a large action space.

Multi-head Puzzle Perception Network

As shown in Fig. 3, the proposed MPPN consists of four main building blocks: A shared encoder to extract common features, a decoder to generate a fragment based on the features, a discriminator head to globally perceive the puzzle and three puzzlet perception heads to perceive the puzzlets.

Encoder and Decoder. The encoder and the decoder form a U-shape network architecture (Liu et al. 2021), where the convolutional layers are placed at the beginning of the encoder and the end of the decoder to focus on low-level image details, and the Swin Transformer blocks replace the convolutional layers in the middle of traditional U-net for feature extraction and pattern recovery at various scales, modeling high-level complex patterns with low computational complexity. Following (Cao et al. 2022; He et al. 2024), each Swin Transformer block contains four Swin Transformer Layers (STLs), and each STL consists of one window multi-head self-attention, and one shifted-window multi-head self-attention, allowing the STLs to capture low-level image fine

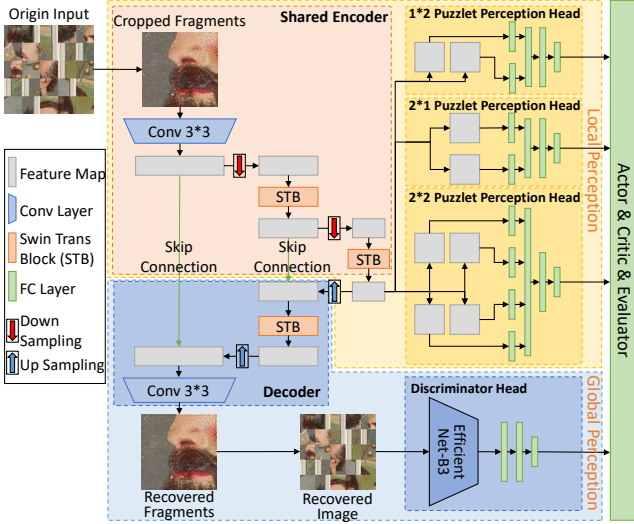


Figure 3: The proposed MPPN with a shared encoder, where a discriminator head perceives global puzzle semantics, and puzzlet perception heads perceive local adjacency relations.

details and high-level image semantics simultaneously (Cao et al. 2022).

Discriminator. The encoder-decoder architecture faces the difficulty of capturing intricate image characteristic and maintaining visual coherence, while a discriminator could help by contrasting the generated fragments with the real ones (Goodfellow et al. 2014). Different from the discriminators applied on puzzle fragments in other models (Bridger, Danon, and Tal 2020; Khoroshiltseva et al. 2022), our discriminator is applied to the recovered puzzle to assess its global semantic features. It is adapted from the EfficientNet-B3 (Tan and Le 2019; Yao et al. 2023) in view of its strong discriminant power, containing 26 MBConv (mobile inverted bottleneck convolution) blocks. Each block consists of a set of convolutional layers and squeeze-and-excitation blocks. The extracted features are fed into three fully connected (FC) layers to aggregate the global image semantics.

Puzzlet Perception Heads. Three puzzlet perception heads are designed to perceive the horizontal and vertical relations between adjacent fragments and the 2×2 puzzlets. As shown in Fig. 3, the heads for horizontal/vertical pairwise relations take the features of two fragments as the input and evaluate their likelihood of being neighbors by aggregating the adjacency information using three fully connected layers. The puzzlet perception head for 2×2 puzzlets takes the features of four fragments as the input to evaluate the likelihood of the puzzlet being correctly assembled or not. All the horizontal neighbors, vertical neighbors, and 2×2 puzzlets are evaluated to derive the local perception of puzzles. They are then aggregated together with the global perception from the discriminator to provide visual clues of the current puzzle.

Aggregation of Visual Clues. Given a puzzle of $M \times N$ fragments, there are $M \times (N - 1)$ pieces of 1×2 puzzlets, $(M - 1) \times N$ pieces of 2×1 puzzlets, and $(M - 1) \times (N - 1)$ pieces of 2×2 puzzlets. To aggregate all the visual clues, we

maximize the evidence E ,

$$E = \lambda_g E_g + E_p, \quad (1)$$

where E_g is global perception from the discriminator, E_p is the perception from the three puzzlet perception heads, and λ_g is a weighting factor. E_p is in turn estimated as,

$$E_p = E_{1 \times 2} + E_{2 \times 1} + E_{2 \times 2}, \quad (2)$$

$$E_{i \times j} = \sum_{x=1}^{M-i+1} \sum_{y=1}^{N-j+1} \lambda_{i \times j}(x, y) E_{i \times j}(x, y), \quad (3)$$

where $i, j \in \{1, 2\}$ are the size of puzzlets, (x, y) denotes the location of puzzlets, $\lambda_{i \times j}(x, y)$ is a weighting factor, and $E_{i \times j}$ is the evidence of puzzlets of size $i \times j$.

Formulation of Markov Decision Process

Puzzle reassembly can be formulated as a combinatorial optimization problem (Song et al. 2023a). Specifically, let Π denote the set of permutations of fragment indices $\{1, 2, 3, \dots, MN\}$ for the puzzle of $M \times N$ pieces. Given an initial permutation $\pi_0 \in \Pi$, the target is to find a mapping function \mathcal{M} so that $\pi_G = \mathcal{M}(\pi_0)$, where π_G is the correct puzzle placement. It is difficult to derive such a mapping \mathcal{M} directly. The target is modified to derive a sequence of mappings $\{\mathcal{M}_1, \mathcal{M}_2, \dots\}$, e.g., a sequence of swapping actions $\pi_{t+1} = \mathcal{M}_t(\pi_t)$ that maximize the evidence E defined in Eq. (1), so that the final permutation reaches π_G . Puzzle reassembly is then formulated as a Markov Decision Process (MDP), represented by the tuple $(\mathbb{S}, \mathbb{A}, \mathcal{R}, P_s)$.

State Space \mathbb{S} , where $s_t \in \mathbb{S}$ is the current state of the puzzle with permutation π_t .

Action Space \mathbb{A} , where $a_t \in \mathbb{A}$ is the action to be applied on the puzzle. Besides common `Swap-2` actions (Song et al. 2023a), two more exploratory actions are designed to expand the action space: `Swap-3` actions that swap three fragments simultaneously, and `Swap-puzzlet` actions that swap two 2×2 puzzlets. The two new types of swapping actions allow the agent to manipulate more puzzle pieces simultaneously, potentially enabling the agent to escape from local optima and converge fast, while `Swap-2` actions allow fine adjustments and refinements of puzzlets, crucial in the final stages to perfectly reassemble the puzzles.

Reward $\mathcal{R}(s_t, a_t)$ for action a_t in state s_t is defined as,

$$\mathcal{R}(s_t, a_t) = \alpha r_t^A + (1 - \alpha) r_t^N + r_t^P + b, \quad (4)$$

where r_t^A rewards the number of fragments being placed at their **Absolute** locations, r_t^N rewards the number of fragment pairs with correct **Neighboring** positions, r_t^P is a large reward for the **Perfect** placement of all fragments and zero otherwise, and b is a small penalty for not correctly reassembling the puzzle at this step. The large reward r_t^P encourages the agent to stop immediately when the puzzle is correctly reassembled, and the small penalty b encourages the agent to converge fast to avoid the lucky success for a long trial. $\alpha \in [0, 1]$ balances the importance of r_t^A and r_t^N . These parameters are empirically determined. More detailed ablations are provided in Supplementary Material.

Transition Probability $P_s(s_{t+1}|s_t)$ is defined as the probability that the agent moves from state s_t to state s_{t+1} .

Algorithm 1: Training EvoRL agent for puzzle reassembly

Input: Number of learning iterations K , maximum number of swaps T , number of evolution iterations Z , action population size S , memory buffer \mathbb{B} , training dataset \mathbb{D}

Output: actor ω_θ , critic ω_φ , and evaluator ω_μ

- 1: **for** $k \leftarrow 1$ **to** K **do**
- 2: Randomly select a puzzle from \mathbb{D}
- 3: **for** $t \leftarrow 1$ **to** T **do**
- 4: Sample a set of actions \hat{a}_t from the actor
- 5: Combine the sampled actions with the historical actions as the initial parents
- 6: **for** $z \leftarrow 1$ **to** Z **do**
- 7: Generate S offspring by crossover and mutation
- 8: Evaluate all offspring by using the evaluator ω_μ and select new parents
- 9: **end for**
- 10: Select the best-performing action a_t from the evolved offspring by using the evaluator ω_μ
- 11: Execute the action a_t onto the puzzle and derive the reward $\mathcal{R}(s_t, a_t)$ as in Eq. (4)
- 12: Store the trajectory $\tau = \{s_t, a_t, \mathcal{R}(s_t, a_t), s_{t+1}\}$ in memory buffer \mathbb{B}
- 13: **end for**
- 14: Sample a batch of transitions τ from \mathbb{B}
- 15: Calculate the state value $\mathcal{V}(s_t)$ as in Eq. (5)
- 16: Compute the advantage function $\mathcal{A}(a_t)$ as Eq. (6)
- 17: Update the actor ω_θ using gradient ascent with the clipped objective function as in Eq. (7)
- 18: Update the critic ω_φ by gradient descent as in Eq. (9)
- 19: Soft update the evaluator ω_μ as in Eq. (10)
- 20: **end for**

EvoRL for Puzzle Reassembly

Existing methods face challenges of handling exponentially increasing search space for L-JPEG problems (Song et al. 2023a; Bridger, Danon, and Tal 2020). In view of its balanced performance between sample efficiency and solution quality (Martinez et al. 2021), an Evolutionary Reinforcement Learning agent is developed for solving L-JPEG problems. As shown in Fig. 2, the agent consists of an actor to recommend the swapping actions from a large action space, a critic to estimate the state value for updating the actor, and an evaluator to assess the evolved transition trajectories and select the leading action from the best trajectory.

The procedure for training the EvoRL agent is summarized in Algo. 1. Given a shuffled puzzle in state s_t , a set of swapping actions $\hat{a}_t \in \mathbb{A}$ are sampled from the action space using the actor policy ω_θ according to the probability $P(\hat{a}_t|s_t)$. The recommended actions \hat{a}_t and the historical actions from the replay buffer are concatenated as the initial parents. Then, the evolutionary operations such as crossover and mutation (Mirjalili 2019) are applied iteratively to produce offsprings, and the evaluator ω_μ assesses the offspring, and selects the action a_t with the highest value as the optimal swapping action. The corresponding reward $\mathcal{R}(s_t, a_t)$ is calculated using Eq. (4). The crossover operation introduces new combinations that explore previously unseen ac-

tion sequences by blending different action sequences, and mutation adds random alterations to action sequences. Both help escape from the local optima by exploring a larger action space. In the initial stage of training, the agent would focus more on broadly exploring the action space by utilizing larger swapping actions to gain a larger reward. When the puzzle is almost perfectly solved, the agent will focus more on exploiting the visual cues of adjacent pieces that are almost correctly placed to better fine-tune the puzzle through smaller swapping actions until it is solved perfectly. After executing the action a_t onto the current puzzle, the puzzle proceeds to the next state s_{t+1} . The trajectory $\tau = \{s_t, a_t, \mathcal{R}(s_t, a_t), s_{t+1}\}$ is stored into the replay buffer \mathbb{B} , where each trajectory is a sequence of states, actions, rewards, and next states. These trajectories are later sampled from \mathbb{B} as the past reassembly experience via the experience replay mechanism.

To guide the training of the actor, the state value $\mathcal{V}(s_t)$ is estimated by the critic as in (Sutton and Barto 2018),

$$\mathcal{V}(s_t) = \mathbb{E}_{a_t, s_{t+1}, \dots} \left[\sum_{l=0}^{\infty} \gamma^l \mathcal{R}(s_{t+l}, a_{t+l}) \right], \quad (5)$$

where $\mathbb{E}_{a_t, s_{t+1}, \dots}$ denotes the expectation over all future feasible states, and γ is the discount factor to avoid greedily maximizing the reward at the current state. For training stability, the advantage function is calculated,

$$\mathcal{A}(s_t, a_t) = \mathcal{R}(s_t, a_t) + \gamma \mathcal{V}(s_{t+1}) - \mathcal{V}(s_t), \quad (6)$$

which evaluates whether the action a_t is better than other possible actions in state s_t . A CLIP function $\mathcal{F}_c(\mathcal{A}(s_t, a_t), \epsilon)$ of the surrogate objective is introduced to constrain the update of $\mathcal{A}(s_t, a_t)$ by a factor of $[1 - \epsilon, 1 + \epsilon]$ as in (Schulman et al. 2017). The actor is then updated by gradient ascent as,

$$L_p = \min[P(a_t|s_t) \cdot \mathcal{A}(s_t, a_t), \mathcal{F}_c(\mathcal{A}(s_t, a_t), \epsilon)] \quad (7)$$

$$\theta \leftarrow \theta + \eta_\theta \cdot \nabla_\theta(L_p), \quad (8)$$

where η_θ is the learning rate. Following (Schulman et al. 2017), the value loss L_v is defined as the squared error of the estimated state value in the trajectory. The critic ω_φ is then updated by gradient descent as,

$$\varphi \leftarrow \varphi - \eta_\varphi \nabla_\varphi \left(\mathcal{V}(s_t) - \sum_{i=t}^T \gamma^{i-t} \mathcal{R}(s_i, a_i) \right)^2, \quad (9)$$

where η_φ is the learning rate. The evaluator assesses all the offspring and selects the optimal swapping action a_t from the offspring. It shares the same architecture as the critic and it is soft-updated by the critic as,

$$\mu \leftarrow \beta * \varphi + (1 - \beta) * \mu, \quad (10)$$

where β is the soft update parameter. The soft-updating mechanism ensures the evaluator to approximate the state value function from the rewards to properly evaluate the offspring, and more importantly, it prevents the evaluator from overfitting caused by the hard update (Wei et al. 2021).

The MPPN is initially trained using a GAN loss. This pre-training phase is critical to align fragments coherently and

discriminate the global puzzle reassembly states. Then, the MPPN is trained with EvoRL agent using the actor and critic losses inherent to the reinforcement learning paradigm, without employing adversarial losses.

During testing, the actor evaluates the visual perception from the MPPN and generates a sequence of swapping actions by using the evolutionary strategy with the assistance from the evaluator till it perfectly assembles the puzzle.

Experimental Results

Experimental Settings

The proposed ERL-MPP is compared with Deepzzle (Pau-mard, Picard, and Tabia 2020), Greedy Search (Paikin and Tal 2015), Tabu Search (Adamczewski, Suh, and Lee 2015), Genetic Algorithm (Sholomon, David, and Ne-tanyahu 2013), Inpaint-GAN (Bridger, Danon, and Tal 2020), PDN-GA (Song et al. 2023b) and SD²RL (Song et al. 2023a) on the following two benchmark datasets.

JPLEG-5 Dataset (Song et al. 2023a) contains 12,000 puzzles, where each puzzle of 534×534 pixels is divided into 5×5 random shuffled pieces with 12-pixel gaps. Following the same experimental setup as in SD²RL (Song et al. 2023a), 9,000 puzzles are selected as the training set, 1,000 as the evaluation set, and 2,000 as the test set.

MIT Dataset (Cho, Avidan, and Freeman 2010) consists of 60 puzzles. Each puzzle is resized and cut into either 7×10 pieces or 10×15 pieces, and each piece has 64×64 pixels. Two types of gaps are applied, 2-pixel gap and 4-pixel gap. Following the same settings as in (Bridger, Danon, and Tal 2020), the models are trained on the DIV2K dataset (Agustsson and Timofte 2017) and tested on this dataset.

For a fair comparison, we adopt the same evaluation metrics as in (Song et al. 2023a), *i.e.*, **Perfect**, **Absolute**, **Horizontal** and **Vertical**, showing the percentage of puzzles that are perfectly reassembled, in their correct absolute positions, in correct horizontal and vertical pairwise relations, respectively. The last two are often combined as **Neighbor** metric.

The hyper-parameters of our ERL-MPP are summarized as follows: the discount rate $\gamma = 0.998$, the learning rate $\eta_\theta = 0.001$ for the actor and $\eta_\phi = 0.001$ for the critic, the soft-update parameter $\beta = 0.001$ for the evaluator, the memory buffer size 200,000, the batch size of 32, the maximum number of iterations $K = 1,000$, the maximum number of swaps $T = 10,000$, the number of evolution iterations $Z = 10$, the action population size $S = 64$, a crossover rate of 90%, a mutation rate of 10%, the weighting factor $\alpha = 0.8$ for absolute reward, $b = 1$ for the penalty and $r_t^P = 1,000$ if the puzzle is perfectly reassembled. The Adam optimizer is used.

Comparison Results on JPLEG-5 Dataset

The comparison results on the JPLEG-5 dataset are summarized in Table 1. It is indeed difficult to reassemble puzzles with large gaps. The two previous best performing methods, PDN-GA and SD²RL, only achieve a **Perfect** rate of 6.1% and 5.1% respectively, and others almost fail to perfectly reassemble any puzzles. The proposed method obtains a **Perfect** rate of 18.6%, significantly better than all the compared

Method	Perf.	Abs.	Hori.	Vert.
Deepzzle (2020)	0.0	21.9	10.9	10.7
Greedy (2015)	0.1	24.1	12.6	12.3
Tabu (2015)	0.0	24.6	12.8	12.8
GA (2013)	0.0	25.1	12.4	12.3
SD ² RL (2023a)	5.1	40.3	26.5	26.2
PDN-GA (2023b)	6.1	44.3	30.8	30.6
Proposed ERL-MPP	18.6	52.7	56.5	57.3

Table 1: ERL-MPP significantly outperforms the compared methods in terms of four metrics on the JPLEG-5 dataset.

Method	Pnt.	Eng.	Art.
Deepzzle (2020)	0.0/15.5	0.0/23.9	0.0/26.3
Greedy (2015)	0.0/16.2	0.0/27.3	0.3/28.7
Tabu (2015)	0.0/16.4	0.0/28.4	0.0/29.0
GA (2013)	0.0/17.0	0.0/28.9	0.0/29.4
SD ² RL (2023a)	0.5/23.8	7.6/48.5	7.2/48.6
PDN-GA (2023b)	0.6/24.5	8.6/54.5	9.1/54.0
Proposed ERL-MPP	3.3/28.5	24.4/62.8	28.1/66.8

Table 2: Comparison results on the JPLEG-5 dataset in terms of **Perfect** (left) and **Absolute** (right) metrics. The proposed ERL-MPP significantly outperforms all compared models.

methods. Compared to the previous best performing method, PDN-GA (Song et al. 2023b), the performance gains are 12.5%, 8.4%, 25.7%, and 26.7% in terms of **Perfect**, **Absolute**, **Horizontal**, and **Vertical** metrics, respectively.

The JPLEG-5 dataset consists of three main types of puzzles: paintings (Pnt.), engravings (Eng.), and artifacts (Art.). The evaluation results on these three types of puzzles are summarized in Table 2. The proposed ERL-MPP significantly outperforms all the compared methods on puzzles of different types using both metrics. Compared to the second best method PDN-GA (Song et al. 2023b), the performance gains are 2.7%, 15.8%, and 19.0%, respectively on painting, engraving, and artifact puzzles in terms of **Perfect** metric, and 4.0%, 8.3%, 12.8% in terms of **Absolute** metric. The painting puzzles are relatively more challenging because of diverse image contents. Thanks to its superior ability to explore and exploit the large action space, the ERL-MPP correctly reassembles many more puzzles than others.

The reassembling results of top-3 performing models are shown in Fig. 4. SD²RL (Song et al. 2023a), and PDN-GA (Song et al. 2023b) often make mistakes for pieces at core areas, while the proposed method could derive the correct reassembly order. The performance gain is mainly attributed to the excellent visual perception capabilities of the MPPN for understanding the image semantics and the capabilities of efficiently and effectively exploring the large action space by the EvoRL agent. For example, in Row 1 of Fig. 4, the up-left fragment corresponding to the lion’s mouth is misplaced by SD²RL, while our ERL-MPP makes use of all visual clues to perfectly reassemble the puzzle, showing the effectiveness of the MPPN in visual perception.

Method	7 × 10 Pieces						10 × 15 Pieces					
	2-pixel gap			4-pixel gap			2-pixel gap			4-pixel gap		
	Perf.	Abs.	Neig.	Perf.	Abs.	Neig.	Perf.	Abs.	Neig.	Perf.	Abs.	Neig.
Greedy (2015)	1.7	42.9	68.4	0.0	12.1	41.4	0.0	42.2	65.9	0.0	11.1	39.7
Tabu (2015)	3.3	49.6	72.3	0.0	20.2	45.8	0.0	48.7	67.3	0.0	12.4	40.3
GA (2013)	3.3	50.1	72.7	0.0	20.1	46.2	0.0	49.3	69.4	0.0	12.5	40.5
Inpaint-GAN (2020)	6.7	86.0	84.6	1.7	50.5	57.1	1.7	66.7	76.3	0.0	35.4	51.3
PDN-GA (2023b)	8.3	86.5	85.3	3.3	52.6	58.8	1.7	67.8	77.8	0.0	37.6	54.8
SD ² RL (2023a)	<u>11.7</u>	<u>86.7</u>	<u>85.4</u>	<u>5.0</u>	<u>53.4</u>	<u>59.2</u>	<u>3.3</u>	<u>68.1</u>	<u>78.0</u>	0.0	<u>38.3</u>	<u>55.3</u>
Proposed ERL-MPP	50.0	90.8	89.2	13.3	59.4	67.3	8.3	70.9	80.4	0.0	40.9	59.7

Table 3: Comparison results on the MIT dataset with 7×10 and 10×15 pieces. The proposed method outperforms all the compared methods in terms of all three evaluation metrics for puzzles of different pieces with different gaps.

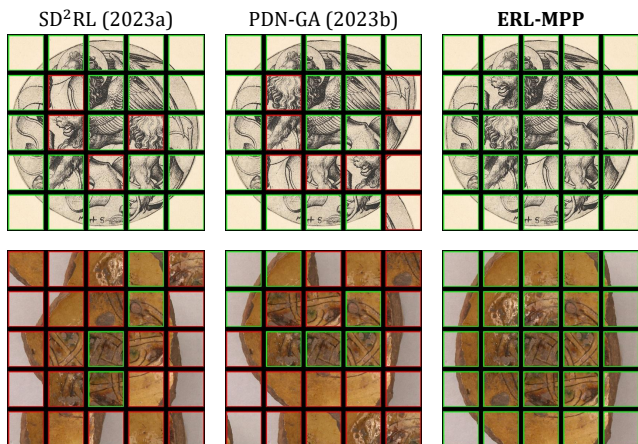


Figure 4: Visualization of the reassembling results on sample puzzles from the JPLEG-5 dataset (Song et al. 2023a).

Comparison Results on MIT Dataset

To evaluate the models on large-scale puzzles, we conduct comparison experiments on the MIT dataset. We adopt four experimental settings, puzzles of 7×10 and 10×15 pieces with 2-pixel and 4-pixel gaps, respectively. The results are summarized in Table 3. The proposed method achieves the best performance among all the compared methods in terms of three evaluation metrics under all four settings. Specifically, for 7×10 puzzles, compared to the second best method SD²RL (Song et al. 2023a), ERL-MPP achieves a **Perfect** rate of 50.0% and 13.3% for 2-pixel gap and 4-pixel gap respectively, approximately as 4 and 3 times high as that of SD²RL. With larger 4-pixel gaps, the performance of all the methods significantly decreases compared to that of 2-pixel gap, showing the difficulty of reassembling large-scale puzzles of large gaps. For 10×15 puzzles, ERL-MPP doubles the perfect reassembled puzzles with 2-pixel gap compared to SD²RL, while none of the methods perfectly assemble any puzzles of 4-pixel gaps. Comparing the results for 7×10 puzzles and 10×15 puzzles, we can observe that the difficulty of puzzle reassembly increases dramatically with the

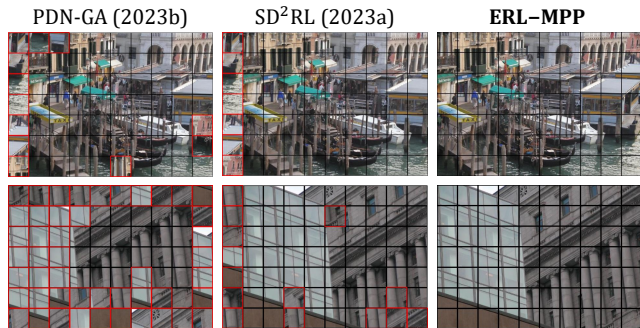


Figure 5: Visualization of the reassembled puzzles of 7×10 pieces with 2-pixel gaps and 4-pixel gaps in the MIT dataset.

number of puzzle pieces, *e.g.*, for 2-pixel gaps, the **Perfect** metric for the proposed ERL-MPP decreases from 50.0% to 8.3% only when the number of puzzle pieces increases from 70 to 150. Dividing a puzzle into too many small pieces with gaps will result in too little semantic information in each fragment and hence no method could work well in this case. Despite all the challenges, the proposed ERL-MPP significantly outperforms all the compared methods in terms of all the three metrics under all the four settings.

Fig. 5 visualizes the final puzzles obtained by top-3 models. Compared to the perfect solutions by the ERL-MPP, most puzzle pieces are well reassembled by SD²RL (Song et al. 2023a) while some incorrect pieces ruin the solution. In contrast, PDN-GA (Song et al. 2023b) does not well reassemble the two puzzles with many pieces wrongly placed. For example, in Row 1 in Fig. 5, SD²RL and PDN-GA are able to align most of the pieces, but they can't perfectly solve the puzzle. The evolutionary strategy helps our method escape from the local optima and perfectly reassemble the puzzle. The results demonstrate the power of ERL-MPP in puzzle reassembly, particularly the importance of evolutionary strategies that efficiently explore more swapping actions and enable the agent to escape from local optima.

Ablation Studies

Ablation Study of Major Components. We ablate two key contributions: MPPN and EvoRL on the JPLEG-5 dataset.

Perception	Agent	Perf.	Abs.	Hori.	Vert.
SDN	DQN	5.1	40.3	26.5	26.2
MPPN	DQN	9.8	46.7	37.8	37.7
SDN	EvoRL	13.5	50.2	43.8	43.8
MPPN	EvoRL	18.6	52.7	56.5	57.3

Table 4: Ablation study of major components of the ERL-MPP on the JPLEG-5 dataset (Song et al. 2023a).

Baseline	2×2	Dis.	Perf.	Abs.	Hori.	Vert.
✓			13.5%	50.2%	43.8%	43.8%
✓	✓		16.4%	51.4%	50.1%	50.3%
✓	✓	✓	18.6%	52.7%	56.5%	57.3%

Table 5: Ablation study of the MPPN on the JPLEG-5 dataset. All heads contribute to the excellent performance.

S-2	S-3	S-P	Perf.	Abs.	Hori.	Vert.
✓			10.0%	46.9%	38.3%	38.1%
✓	✓		14.7%	49.7%	46.5%	46.8%
✓	✓	✓	18.6%	52.7%	56.5%	57.3%

Table 6: Ablation of action space on the JPLEG-5 dataset.

SD²RL (Song et al. 2023a) is selected as the baseline, which utilizes the Siamese Discriminant Network (SDN) for puzzle perception and the Deep Q-Network (DQN) for puzzle reassembly. Table 4 presents the ablation results by gradually replacing SDN and DQN by MPPN and EvoRL. Both MPPN and EvoRL bring significant gains. The performance gain of the MPPN over the SDN is 5.1%, 2.5%, 12.7%, and 13.5% in terms of the **Perfect**, **Absolute**, **Horizontal**, and **Vertical** metrics, respectively when using the EvoRL agent, demonstrating the power of the shared-encoder and multi-head perception in extracting both local and global puzzle information. Compared to the DQN, the EvoRL achieves the performance gain of 8.8%, 6.0%, 18.7%, and 19.6% respectively when using the MPPN for visual perception, demonstrating the effectiveness of the EvoRL agent in exploring and exploiting the large action space. The ablation results confirm the contributions of both core components to the excellent performance.

Ablation Study of MPPN. We ablate the three network heads of the proposed MPPN on the JPLEG-5 dataset as shown in Table 5. The MPPN with 1×2 and 2×1 heads serves as the baseline. By adding the 2×2 head, the performance increases in terms of all four evaluation metrics. Furthermore, by adding the discriminator head, the proposed MPPN achieves a further performance gain of 2.2%, 1.3%, 6.4%, and 7.0% respectively in terms of **Perfect**, **Absolute**, **Horizontal**, and **Vertical** metrics. The ablation results show that all the heads contribute significantly to the final results.

Ablation Study of Action Space. The action space plays a key role in solving the puzzles. We ablate it to justify our design. Swap-2 (S-2) is selected as the baseline, which has been used exclusively (Song et al. 2023a). Swap-3 (S-3)

Method	Training	Test	Model Size
Deepzlle (2020)	5.6h	7.7s	15.8M
Greedy (2020)	10.2h	5.6s	31.5M
Tabu (2015)	10.2h	9.8s	31.5M
GA (2013)	10.2h	14.3s	31.5M
SD ² RL (2023a)	30.7h	44.7s	44.2M
PDN-GA (2023b)	19.7h	32.5s	46.8M
Proposed ERL-MPP	36.2h	53.9s	102.1M

Table 7: Comparison of execution time and model size on the JPLEG-5 dataset.

and Swap-puzzlet (S-P) are added for ablation. The experimental results are summarized in Table 6. It is obvious that by extending the action space from Swap-2 actions to Swap-3 and Swap-puzzlet actions, all the evaluation metrics significantly improve, demonstrating the power of the new action space.

Time Complexity Analysis

The execution time and model size for all the compared methods on the JPLEG-5 dataset are summarized in Table 7. The training time is obtained over all the images in the training dataset, while the test time is averaged per test image. Compared with other methods, the RL methods, *i.e.*, SD²RL (Song et al. 2023a) and the proposed ERL-MPP, require more time for the agents to train and infer the correct reassembly sequence, as the agent needs to extract the visual features through the deep convolution networks from the puzzle in every reassembly step. The proposed ERL-MPP is larger than all the compared models, as it contains an Evolutionary Reinforcement Learning agent with Actor, Critic, and Evaluator network, and a more complex visual perception module MPPN for extracting both global and local features. It exhibits a significant performance gain at a price of longer execution time and larger model size.

Conclusion

The proposed ERL-MPP tackles the challenges of solving L-JPEG problems. Specifically, the proposed MPPN visually perceives both the global semantics through the discriminator head and the local puzzle status through the three puzzlet perception heads, which provides informative visual clues for the EvoRL agent. To tackle the challenges of reassembling a large number of puzzle fragments with gaps, the proposed EvoRL agent selects a sequence of swapping actions based on both the current visual perception and past reassembly experience embedded in the agent. The novel actor-critic-evaluator architecture in conjunction with the evolutionary strategy demonstrates superior performance and efficiently explores and exploits the large action space featured with complex landscapes. Extensive experiments on puzzles with large gaps in the JPLEG-5 dataset and puzzles with large numbers of puzzle pieces with gaps in the MIT dataset show that the proposed ERL-MPP significantly outperforms all the compared methods in terms of all evaluation metrics under various scale and gap settings.

Acknowledgments

This work was supported by the Ningbo Municipal Bureau of Science and Technology under Grant 2022J171, 2022Z173, 2022Z217, 2023Z138, 2023Z237 and 2024Z110.

References

- Adamczewski, K.; Suh, Y.; and Lee, K. M. 2015. Discrete tabu search for graph matching. In *Proc. IEEE Int. Conf. Comput. Vis.*, 109–117.
- Agustsson, E.; and Timofte, R. 2017. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit. Workshops*, 126–135.
- Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; and Bharath, A. A. 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*.
- Bridger, D.; Danon, D.; and Tal, A. 2020. Solving Jigsaw puzzles with eroded boundaries. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 3526–3535.
- Cao, H.; Wang, Y.; Chen, J.; Jiang, D.; Zhang, X.; Tian, Q.; and Wang, M. 2022. Swin-Unet: Unet-like pure transformer for medical image segmentation. In *Eur. Conf. Comput. Vis.*, 205–218. Springer.
- Cho, T. S.; Avidan, S.; and Freeman, W. T. 2010. A probabilistic image Jigsaw puzzle solver. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 183–190.
- Derech, N.; Tal, A.; and Shimshoni, I. 2021. Solving archaeological puzzles. *Pattern Recognit.*, 119: 108065.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised visual representation learning by context prediction. In *Proc. IEEE Int. Conf. Comput. Vis.*, 1422–1430.
- Du, R.; Chang, D.; Bhunia, A.; Xie, J.; Ma, Z.; Song, Y.; and Guo, J. 2020. Fine-grained visual classification via progressive multi-granularity training of Jigsaw patches. In *Eur. Conf. Comput. Vis.*, 153–168. Springer.
- Gallagher, A. C. 2012. Jigsaw puzzles with pieces of unknown orientation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 382–389.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Adv. Neural Inf. Process. Syst.*, volume 27.
- Harel, P.; Shahar, O.; and Ben-Shahar, O. 2024. Pictorial and apictorial polygonal Jigsaw puzzles from arbitrary number of crossing cuts. *Int. J. Comput. Vis.*, 1573–1405.
- He, W.; Ren, J.; Bai, R.; and Jiang, X. 2024. Hierarchical Perceptual and Predictive Analogy-Inference Network for Abstract Visual Reasoning. In *Proc. ACM Int. Conf. Multimedia*, 4841–4850.
- Khoroshiltseva, M.; Traviglia, A.; Pelillo, M.; and Vascon, S. 2022. Relaxation labeling meets GANs: Solving Jigsaw puzzles with missing borders. In *Proc. Int. Conf. Image Anal. Process.*, 27–38.
- Li, C.; He, Y.; Ren, J.; Bai, R.; Zhao, Y.; Yu, H.; and Jiang, X. 2024. Regression Residual Reasoning with Pseudo-labeled Contrastive Learning for Uncovering Multiple Complex Compositional Relations. In *Int. Jt. Conf. Artif. Intell.*, volume 4, 3466–3474.
- Li, R.; Liu, S.; Wang, G.; Liu, G.; and Zeng, B. 2022. JigsawGAN: Auxiliary learning for solving Jigsaw puzzles with generative adversarial networks. *IEEE Trans. Image Process.*, 31: 513–524.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. IEEE Int. Conf. Comput. Vis.*, 10012–10022.
- Ma, C.; Rao, Y.; Lu, J.; and Zhou, J. 2021. Structure-preserving image super-resolution. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(11): 7898–7911.
- Martinez, A. D.; Del Ser, J.; Osaba, E.; and Herrera, F. 2021. Adaptive multifactorial evolutionary optimization for multi-task reinforcement learning. *IEEE Trans. Evolut. Comput.*, 26(2): 233–247.
- Mirjalili, S. 2019. Genetic algorithm. In *Evolutionary algorithms and neural networks*, 43–55. Berlin/Heidelberg, Germany: Springer.
- Nada, A. R.; and Jan, N. 2015. A survey of computer methods in reconstruction of 3D archaeological pottery objects. *J. Adv. Res.*, 3(3): 712–714.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving Jigsaw puzzles. In *Proc. Eur. Conf. Comput. Vis.*, 69–84.
- Paikin, G.; and Tal, A. 2015. Solving multiple square Jigsaw puzzles with missing pieces. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 4832–4839.
- Paumard, M. M.; Picard, D.; and Tabia, H. 2018. Image reassembly combining deep learning and shortest path problem. In *Proc. Eur. Conf. Comput. Vis.*, 153–167.
- Paumard, M. M.; Picard, D.; and Tabia, H. 2020. Deepzle: Solving visual Jigsaw puzzles with deep learning and shortest path optimization. *IEEE Trans. Image Process.*, 29: 3569–3581.
- Ren, B.; Liu, Y.; Song, Y.; Bi, W.; Cucchiara, R.; Sebe, N.; and Wang, W. 2023. Masked Jigsaw puzzle: A versatile position embedding for vision transformers. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 20382–20391.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sholomon, D.; David, O.; and Netanyahu, N. S. 2013. A genetic algorithm-based solver for very large Jigsaw puzzles. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 1767–1774.
- Son, K.; Hays, J.; and Cooper, D. B. 2014. Solving square Jigsaw puzzles with loop constraints. In *Eur. Conf. Comput. Vis.*, 32–46.
- Son, K.; Hays, J.; and Cooper, D. B. 2019. Solving square Jigsaw puzzle by hierarchical loop constraints. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(9): 2222–2235.

Song, X.; Jin, J.; Yao, C.; Wang, S.; Ren, J.; and Bai, R. 2023a. Siamese-discriminant deep reinforcement learning for solving Jigsaw puzzles with large eroded gaps. *Proc. AAAI Conf. Artif. Intell.*, 37(2): 2303–2311.

Song, X.; Yang, X.; Ren, J.; Bai, R.; and Jiang, X. 2023b. Solving Jigsaw puzzle of large eroded gaps using puzzlet discriminant network. In *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. Cambridge, MA: MIT press.

Tan, M.; and Le, Q. 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proc. Int. Conf. Mach. Learn.*, 6105–6114.

Wei, Z.; Quan, Z.; Wu, J.; Li, Y.; Pou, J.; and Zhong, H. 2021. Deep deterministic policy gradient-DRL enabled multiphysics-constrained fast charging of lithium-ion battery. *IEEE Trans. Ind. Electron.*, 69(3): 2588–2598.

Yan, F.; Zheng, Y.; Cong, J.; Liu, L.; Tao, D.; and Hou, S. 2021. Solving Jigsaw puzzles via nonconvex quadratic programming with the projected power method. *IEEE Trans. Multimedia*, 23: 2310–2320.

Yang, X.; Wang, Y.; Chen, K.; Xu, Y.; and Tian, Y. 2022. Fine-grained object classification via self-supervised pose alignment. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 7399–7408.

Yao, C.; Ren, J.; Bai, R.; Du, H.; Liu, J.; and Jiang, X. 2023. Mask attack detection using vascular-weighted motion-robust rPPG signals. *IEEE Trans. Inf. Forensics Secur.*

Zhang, J.; Ren, J.; Zhang, Q.; Liu, J.; and Jiang, X. 2023. Spatial context-aware object-attentional network for multi-label image classification. *IEEE Trans. Image Process.*, 32: 3000–3012.

Zhang, J.; Wang, X.; Yao, C.; Ren, J.; and Jiang, X. 2024a. Visual-linguistic Cross-domain Feature Learning with Group Attention and Gamma-correct Gated Fusion for Extracting Commonsense Knowledge. In *Proc. ACM Int. Conf. Multimedia*, 4650–4659.

Zhang, J.; Yang, X.; He, W.; Ren, J.; Zhang, Q.; Zhao, Y.; Bai, R.; He, X.; and Liu, J. 2024b. Scale Optimization Using Evolutionary Reinforcement Learning for Object Detection on Drone Imagery. *Proc. AAAI Conf. Artif. Intell.*, 38(1): 410–418.

Zhang, K.; Yu, W.; Manhein, M.; Waggenspack, W.; and Li, X. 2015. 3D fragment reassembly using integrated template guidance and fracture-region matching. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 2138–2146.