Interpreting Graph Drawing with Multi-Agent Reinforcement Learning

Ilkin Safarli* Youjia Zhou† Bei Wang‡

Scientific Computing and Imaging Institute, University of Utah

ABSTRACT

Applying machine learning techniques to graph drawing has become an emergent area of research in visualization. In this paper, we interpret graph drawing as a multi-agent reinforcement learning (MARL) problem. We first demonstrate that a large number of classic graph drawing algorithms, including force-directed layouts and stress majorization, can be interpreted within the framework of MARL. Using this interpretation, a node in the graph is assigned to an agent with a reward function. Via multi-agent reward maximization, we obtain an aesthetically pleasing graph layout that is comparable to the outputs of classic algorithms. The main strength of a MARL framework for graph drawing is that it not only unifies a number of classic drawing algorithms in a general formulation, but also supports the creation of novel graph drawing algorithms by introducing a diverse set of reward functions.

1 INTRODUCTION

Graphs are widely used to model complex relational data such as computer networks, brain connectomics, protein-protein interactions, and communication patterns within social media. To visualize graphs, various techniques have been developed in the graph drawing and information visualization literature (see e.g., [18, 19, 48] for surveys).

Node-link diagrams are perhaps one of the most popular and intuitive graph visualization methods, where nodes are represented as points and edges as lines. However, drawing a graph with an arbitrary node-link diagram, on its own, does not necessarily lead to insight into its underlying data. The layout of nodes and edges strongly influences how a graph and its underlying data are understood. Therefore, an important point in graph drawing is "not simply drawing a graph but how the graph is drawn" [18]. Depending on the structure of the data and the nature of the application domain, various styles of layouts, such as force-directed, circular, orthogonal, and hierarchical, have been developed. Since no universal layout algorithm exist that can produces good results for all types of graphs, we ask a different question instead: *Can we put a large number of classic graph drawing algorithms in a unifying framework?*

Even though many layout algorithms have been developed since Tutte [50] first introduced his barycenter method in 1960, the graph drawing problem remains an active area of research. Recently, applying machine learning techniques to graph drawing has started to gain traction (e.g. [30, 31, 52]). Many of these approaches rely on classic graph drawing algorithms to generate a training dataset, and then utilize machine learning models to learn from the training dataset. However, the training process can be time-consuming, its performance relies heavily on the similarity between the training and testing datasets, and the models typically do not generalize for different graphs and different layouts (e.g. [31]). We take a different

---

*e-mail: ilkin@sci.utah.edu

†e-mail: zhou325@sci.utah.edu

‡e-mail: beiwang@sci.utah.edu

perspective: *What are the different ways classic graph drawing algorithms can be reinvented or transformed using machine learning?*

In this paper, instead of a deep-learning-flavored approach (e.g. [31, 52]) that relies on training datasets, we take a broader view of learning and focus on the paradigm of reinforcement learning, where we learn to "choose the correct actions based on the outcomes of previous actions in similar situations" [7]. Specifically, we interpret the graph drawing problem as a multi-agent reinforcement learning (MARL) problem.

We first demonstrate that a large number of classic graph drawing algorithms, including force-directed layouts, stress majorization, and incremental layouts, can be interpreted within the framework of MARL. We achieve this interpretation by assigning agents to nodes with reward functions that are derived from classic layout algorithms. An aesthetically pleasing layout that is comparable to the classic algorithm is obtained via multi-agent reward maximization. We then propose a set of novel graph drawing algorithms by introducing a diverse set of reward functions. To the best of our knowledge, this is the first time reinforcement learning is utilized in graph drawing.

We investigate the effectiveness of the proposed framework through both qualitative and quantitative evaluations based on established metrics. We develop a tool call MARLL (pronounce as "Mar Roll", which stands for Multi-Agent Reinforcement Learning Layouts) to demonstrate the utility of our MARL approach. MARLL is available via Github[1]. It implements a set of classic layout algorithms and their MARL counterparts and allows the comparison of their resulting layouts across synthetic and real world datasets.

In summary, our contributions include:

- A novel, unifying framework that interprets a large number of classic graph drawing algorithms with MARL.
- A method for recombining existing and deriving new layout algorithms based on a diverse set of reward functions.
- A demonstration of the effectiveness of our method via qualitative and quantitative evaluations.
- An open-source tool that implements and demonstrates various classic and MARL layouts that facilitates visual comparisons.

2 RELATED WORK

A number of studies have applied machine learning to graph drawing, see [51] for a survey. We classify these studies based on the models they employ, namely, evolutionary learning, neural networks/deep learning, and other approaches. We also give a brief overview of related work in reinforcement learning. To the best of our knowledge, this is the first time reinforcement learning is used to study graph drawing algorithms.

Evolutionary learning. Evolutionary algorithms are one of the early machine learning approaches for graph drawing. The positions of vertices in a graph layout represent the genotypes of individuals. The fitness function is a weighted sum of functions that measures how well a drawing accommodates certain aesthetic criteria [41]. Based on a fitness function, best candidate layouts are selected at each step. Crossover and mutation processes are then applied to the candidates to generate a new layout until the layout converges.

---

[1] https://github.com/kinimesi/marl-layout-demo

The approaches that use evolutionary learning can be further classified into two categories: independent from or dependent on user interactions. Approaches in the first category automatically compute the layout of a given graph without any user interaction (e.g., [28, 46, 49]). The fitness function encodes various aesthetic criteria, such as minimizing node overlaps (e.g., [27, 28]), minimizing the number of edge crossings (e.g., [28, 49]), and maximizing the uniformity of edge lengths (e.g., [27, 49]), see [41] for a survey.

Many approaches in the second category follow a similar pipeline via evolutionary learning with user interactions (e.g. [1, 32, 41, 46]): users are first asked to rate a set of graph layouts; a fitness function that reflects their preferences is learned based on user feedback, and then a layout is produced from the fitness function. In the work by Masui [32], users provided the system with pairs of good and bad layout examples, and the system inferred the fitness function from these examples, which then evolved via genetic programming. Rosete-Suarez *et al.* [41] built a system that acquires user preferences either from user evaluations or from user actions.

Neural networks for graph drawing. Cimikowski and Shope [8] presented a parallel neural network algorithm for minimizing the number of edge crossing of nonplanar graphs in a linear embedding. Each edge in the graph is associated with an "up" or "down" neuron, which represents an embedding of the edge in the upper or lower half-plane [8]. Two kinds of forces – excitatory and inhibitory forces – push the neurons iteratively to a state that minimizes the number of edge crossings [51]. Their system does not build a neural network for learning purpose, but rather it models the graph structure as a neural network coupled with an energy system [51].

Wang *et al.* [52] proposed a graph-LSTM-based model to learn and predict the drawing of an input graph. LSTM (long short-term memory) is a popular version of recurrent neural networks (RNNs). Using a set of layout examples as the training dataset, a LSTM model was trained to learn the characteristics of the training dataset and its corresponding algorithm-specific parameters for determining desirable graph drawings. In particular, direct connections were added between different LSTM cells to explicitly model the topological structure of input graphs. Given a new input graph, the trained model was then used to generate a graph drawing with visual properties similar to those of the training dataset.

Kwon and Ma [31] introduced a deep generative model based on VAEs (variational autoencoders) [23] that systematically visualizes a graph in diverse layouts. They present an encoder-decoder architecture to learn a VAE from a collection of example layouts, where the encoder represents training examples in a latent space and the decoder generates layouts from the latent space. A what-you-see-is-what-you-get (WYSIWYG) interface was constructed by mapping the generated samples onto the 2D latent space. Such an interface allowed users to navigate and generate among various representative layouts without adjusting parameters of layout methods [31]. However, although a trained model could generate different layouts for the same input graph, a new model needs to be trained for each new input. Generalizing such a model to an unseen graph remains challenging.

Other machine learning approaches for graph drawing. Kwon *et al.* [30] proposed a machine learning approach to the graph drawing problem that gives the user an idea of "what a graph would look like" in a particular layout. Given an input graph, they did not compute its layout directly. Instead, they found the most topologically similar graph to the input graph using graphlet frequency, and then displayed its pre-computed layout. Their method can be quite useful, especially when an input graph is large and the user is interested only in an approximated layout. However, their method needs a training dataset that includes the layouts of many graphs and their quality metrics, which require a considerable amount of time to generate.

Reinforcement learning. A foundational idea underlying theories of learning and intelligence is *learning from interaction*, and reinforcement learning (RL) is its corresponding computational approach [47]. There have been many recent advances in applying RL to well-known sequential decision-making problems; see [47] for an introduction and [21, 55] for surveys. Successful applications of RL include the AlphaGo system, which defeated a professional human Go player [44, 45], robotics control [25], playing card games [4], and autonomous driving [43]. RL systems are divided into two categories, single-agent RL and multi-agent RL. As the name suggests, they differ by the number of agents operating in the system and the associated RL algorithms. We give some technical background on these systems in Sect. 3. To the best of our knowledge, RL has not been applied in graph drawing as presented in this paper.

3 TECHNICAL BACKGROUND

3.1 Graph Drawing

An unweighted, undirected graph $G = (V, E)$ consists of a node set V and an edge set $E \subseteq V \times V$ connecting pairs of nodes. Let $|V| = n$ and $|E| = m$. The problem of graph drawing is to find a graph layout – a set of coordinates p_v for each node $v \in V$ – such that the drawing is aesthetically appealing [20]. In this paper, we consider 2D graph drawings where $p_v \in \mathbb{R}^2$. We also assume that edges are drawn as straight lines, and nodes have uniform sizes and circular shapes.

Force-directed layout. Force-directed graph layout is a class of algorithms that is widely used for simple undirected graphs [18, 26]. The graph is modeled as a physical system where nodes are attracted and repelled based on various force formulations.

Eades's spring-embedder inspired many force-directed layout algorithms [13]. He proposed a mechanical model in which the nodes are represented with steel balls and the edges with springs. The nodes are placed with a random configuration and released. The force on each node in the system is the sum of repulsive forces from all nodes and attractive forces from adjacent nodes. A layout is obtained once the system reaches equilibrium.

FR layout. Fruchterman and Reingold presented a modification of the spring-embedder model of Eades with increased speed and simplicity [15] (referred to as a FR layout in this paper). Their principles for graph drawing remain applicable even today: "nodes connected by an edge should be drawn near each other" and "nodes should not be drawn too close to each other" [15]. Fruchterman and Reingold were concerned with aesthetic criteria, including even node distribution, minimized edge crossings, uniform edge lengths, inherent symmetry, and conforming the drawing to a given frame. In their model, the nodes are represented as atomic particles or celestial bodies that "exert attractive and repulsive forces on one another" to induce movement [15].

First, the attractive forces and repulsive forces between a pair of nodes u and v are calculated with

$$f_a(u, v) = \frac{d(u, v)^2}{k}, \tag{1}$$

$$f_r(u, v) = \frac{-k^2}{d(u, v)}, \tag{2}$$

where $d(u, v) := ||p_u - p_v||$ is the distance. Equation (1) is computed for $(u, v) \in E$ and Equation (2) is for all pairs of nodes. k is the optimal distance between nodes; it is calculated as $k = C\sqrt{\frac{\text{area}}{|V|}}$, where area is the size of the drawing canvas and the constant C is found experimentally. Then, the notion of temperature is added to control the displacement of nodes, so that as the layout becomes better, the adjustments of the nodes become smaller with respect to the decreasing in temperature [26].

DGC layout. Dogrusoz *et al.* [12] proposed another spring-embedder variant for laying out general graphs with non-uniform

node sizes and compound nodes (referred to as a DCG layout here). Their force model also consists of both attractive and repulsive forces, calculated as

$$f_a(u,v) = \frac{(\lambda - d(u,v))^2}{\zeta}, \quad (3)$$

$$f_r(u,v) = \frac{\mu}{d(u,v)^2} \quad (4)$$

where λ is the ideal edge length, ζ is the elastic constant of the edge, and μ is the repulsion constant. Similar to Fruchterman and Reingold, they also limited the maximum node displacement based on a temperature cooling schema.

Stress majorization. Gansner *et al.* [17] proposed to perform energy optimization with stress majorization. Their stress function is defined as

$$E = \sum_{u<v} w_{u,v}(d(u,v) - d'(u,v))^2 \quad (5)$$

where $d'(u,v)$ is the graph-theoretic distance between u and v, and $w_{u,v}$ is the normalization constant, which equals $d'(u,v)^{-2}$. This stress function can be globally optimized via majorization (a method from multidimensional scaling), which is guaranteed to converge [26]. This approach can also improve the layout quality and running time in practice [17]. Zheng *et al.* [56] used stochastic gradient descent to minimize the same stress function.

3.2 Multi-Agent Reinforcement Learning

We review the technical formulations for finite single-agent reinforcement learning (SARL) and its multi-agent counterpart following the notations from the survey of Zhang *et al.* [55] with minor modifications based on the work of Busoniu *et al.* [5].

SARL. In a dynamic environment, an RL agent is an algorithmic component that learns by interacting with the environment via trials and errors [21]. An agent senses the state of the environment and takes an action; its purpose is to reach its goal by taking actions that maximize its associated reward [47]. Such a process can be modeled as a Markov decision process (MDP).

Formally, a SARL system can be formulated as an MDP via a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ [55]:

- $\mathcal{S}$ is the *state space* that contains states of the environment.
- $\mathcal{A}$ is the *action space* consisting of actions available to the agent.
- $\mathcal{P}(s,a,s')$ is the *transition probability* from any state $s \in \mathcal{S}$ to any state $s' \in \mathcal{S}$ for any given action $a \in \mathcal{A}$.
- $\mathcal{R}(s,a,s')$ is the *reward function* that determines the immediate reward received by the agent after transitioning from s to s' with action a.
- γ is the *discount factor* that balances between the instantaneous and future rewards.

The behavior of the agent is described by its policy π, which is a mapping from the state space $\mathcal{S}$ to the distribution over the action space $\mathcal{A}$ that describes how the agent chooses its actions given the state [5]. $\mathcal{P}$, $\mathcal{R}$, and π can be deterministic or stochastic. In a deterministic model, the next state, the reward, and the policy are completely determined by the current state-action pair. In the stochastic model, they are drawn from distributions. We describe the stochastic model here.

The goal is to maximize the expected *discounted accumulated reward* by finding a policy π [55],

$$\mathbb{E}\left[\sum_{t\geq 0} \gamma^t \mathcal{R}(s_t, a_t, s_{t+1}) \,\middle|\, a_t \sim \pi(s_t), s_0\right], \quad (6)$$

where $a_t \sim \pi(s_t)$ is the action drawn from the policy, and the expectation is taken over $s_{t+1} \sim \mathcal{P}(s_t, a_t, s_{t+1})$.

One way to achieve this maximization is by computing an optimal *Q-function* (i.e., action-value function) for a state-action pair (s,a) by capturing its expected return given the policy π [5]:

$$Q_\pi(s,a) = \mathbb{E}\left[\sum_{t\geq 0} \gamma^t \mathcal{R}(s_t, a_t, s_{t+1}) \,\middle|\, a_t \sim \pi(s_t), a_0 = a, s_0 = s\right]. \quad (7)$$

Various RL methods have been developed to find a good estimate of the optimal Q-function. The (approximate) optimal policy can then be obtained by taking the ε-greedy action of the Q-function estimate. One of the most popular methods for estimating the optimal Q-function is the Q-learning algorithm [53]. For a given agent that takes action a from state s to state s' and gets a reward $r = \mathcal{R}(s,a,s')$, the algorithm adjusts the estimated value of Q-function as

$$\hat{Q}_\pi(s,a) \leftarrow (1-\alpha)\hat{Q}_\pi(s,a) + \alpha\left[r + \gamma\max_{a'}\hat{Q}_\pi(s',a')\right] \quad (8)$$

where α is the learning rate, and γ is the discount factor. The approximate optimal policy is obtained by taking the greedy action of the Q-function estimate [55]. That is, an agent picks for every state s the action a with the highest value of the Q-function [5].

MARL. Using a generalized version of the MDP, a MARL system can be formulated as a Markov game (MG). It is specified by a tuple $(\mathcal{N}, \mathcal{S}, \{\mathcal{A}^i\}_{i\in\mathcal{N}}, \mathcal{P}, \{\mathcal{R}^i\}_{i\in\mathcal{N}}, \gamma)$ [55]. $\mathcal{N} = [1, \cdots, N]$ denotes the set of $N > 1$ agents. $\mathcal{A}^i$ and $\mathcal{R}^i$ are the state space and the reward function of agent i, respectively.

MARL systems are mainly divided into three main categories: fully cooperative, fully competitive, and mixed systems [5, 55]. We implement the fully cooperative version in this paper, as the name suggests, all the agents share the same reward function and Q-function. It would be interesting to explore fully competitive and mixed systems in the future and observe their effect on graph layouts.

4 METHOD

Our proposed approach models and thus interprets a number of graph drawing algorithms with MARL. To the best of knowledge, this is the first time RL is utilized in graph drawing.

4.1 Model Graph Drawings with MARL

We interpret the graph drawing problem with MARL by assigning an agent to each node. During each step of an iterative process, the agents can take actions and move toward the direction within the environment that maximizes their reward function. A graph layout is obtained via reward maximization.

States and actions. In our model, we divide the state space of a given agent into a 3×3 grid as shown in Fig. 1. There are in total nine states. An agent can move from any one of the nine total states to any other state by taking the relevant action. The action space consists of nine actions accordingly: stay (do not move), move north, south, east, west, northeast, northwest, southeast, and southwest, respectively. All actions are possible from any state.

Reward function. The reward function is the feedback mechanism for an agent that defines the consequences of the agent's actions and facilitates learning by interaction [47]. In many force-directed and energy-based graph drawing algorithms, the equilibrium state is obtained by minimizing the energy or the stress of the system. In our model, we formulate the energy or the stress via a reward function. Suppose we assign an agent to a node in the graph. The difference of the force or energy (stress) on each node before and after the action of an agent is defined to be the reward function for a state-action
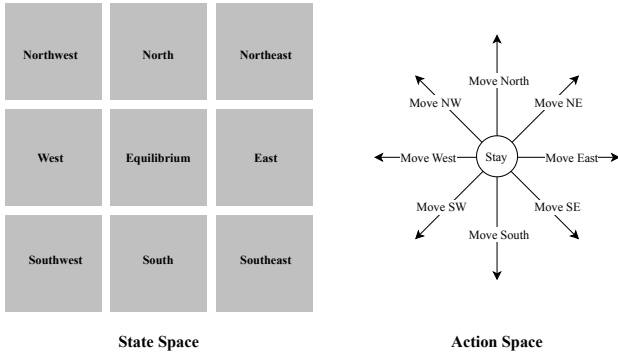
Figure 1: The state and action space of a given agent. An agent can take any action from any state.

pair. Specifically, for force-based formulation, the reward function for an agent v at time (i.e., iteration) t is defined as

$$\mathcal{R}^v(t) = \|F^v(t-1)\| - \|F^v(t)\|, \qquad (9)$$

where $F^v(t)$ is the total force on node v at time t, $||\cdot||$ is the vector norm. For energy-based formulation, the reward function is

$$\mathcal{R}^v(t) = E^v(t-1) - E^v(t). \qquad (10)$$

The goal of any agent is to maximize its reward by moving the corresponding node around according to a policy that is learned by interacting with the environment. The formulations in Equation (9) and Equation (10) provide a general framework to interpret a number of graph drawing algorithms that are based on physical simulations. In Sect. 4.2, we discuss how different algorithms can be interpreted with this formulation.

Finding optimal policy. In MARL, Q-learning is an iterative approximation procedure that estimates the optimal value function [5, 53]. For a given agent, the Q-function is given in Equation (8). At each iteration, the action of an agent is selected based on an ε-greedy policy [5] (reviewed in Sect. 3). If the associated reward of an action is positive, then the action is taken by moving the node to its new position. If the reward is negative, then an action is accepted with a small probability to prevent the algorithm from becoming stuck at a local minimum (similar to a Metropolis–Hastings algorithm). The learning process continues until the system converges.

Convergence criteria. We consider the convergence of a Q-learning algorithm as the convergence of its underlying graph layout. We demonstrate via experiments in Sect. 7 that at convergence, we arrive at an aesthetically pleasing graph layout. Our model has four criteria for convergence. The algorithm terminates whenever one of these criteria is reached. The first criterion M is simply the total number of (allowable) iterations.

Second, if the average node displacement is less than a predefined threshold, then the algorithm terminates. The *average node displacement A* at time t is given by $A(t) = \frac{\sum_{v \in V} \|p_v(t) - p_v(t-1)\|}{n}$, where n is the number of nodes, and $p_v(t)$ is the position of node v at time t. $\|\cdot\|$ denotes the Euclidean norm.

Third, if the node displacement rate is smaller than a predefined threshold, then the algorithm terminates. The *displacement rate ΔA* at time t is calculated by $\Delta A(t) = \|A(t) - A(t-1)\|$.

Finally, we utilize the stress ratio as a termination criterion according to Gansner *et al.* [17], when the reward function is based on the energy of the system. That is, the algorithm terminates if the stress ratio falls below a predefined threshold. The *stress ratio δE* at time t is given as $\delta E(t) = \frac{E(t) - E(t-1)}{E(t)}$, where $E(t)$ is the total energy of the system at time t.

4.2 Interpret Classical Layouts With MARL

In Equation (9) and Equation (10), we present a general formula for the reward function. We now discuss how to interpret a number of classical graph drawing algorithms with MARL based on such a formulation. For the FR and DCG layouts, the reward function for agent v at time t follows Equation (9). For the stress majorization, the reward function is given by Equation (10).

FR layout. The layout algorithm proposed by Fruchterman and Reingold [15] can be adjusted to our model by rewriting the total force on node v at time t as $F^v(t) = \sum_{(u,v) \in E} f_a^t(u,v) + \sum_{u \in V} f_r^t(u,v)$, for attractive forces f_a^t and repulsive forces f_r^t at time t. $F^v(t)$ can be written as

$$F^v(t) = \sum_{(u,v) \in E} \frac{d_t(u,v)^2}{k} + \sum_{u \in V} \frac{-k^2}{d_t(u,v)}, \qquad (11)$$

where $d_t(u,v)$ is the distance between u and v in the layout at time t.

DGC layout. Similarly, for the algorithm proposed by Dogrusoz *et al.* [12], the total force on node v at time t is written as

$$F^v(t) = \sum_{(u,v) \in E} \frac{(\lambda - d_t(u,v))^2}{\zeta} + \sum_{u \in V} \frac{\mu}{d_t(u,v)^2}. \qquad (12)$$

Stress majorization. The optimization of the stress function used by Gansher *et al.* [17] can be formulated in two ways via MARL: optimizing the local stress on each agent, or optimizing the global stress via the action of an agent. To optimize the local stress on each agent, we define the *local stress* for agent v at time t as

$$E^v(t) = \sum_{u \in N(v,p)} w_{u,v}(d_t(u,v) - d'_t(u,v))^2, \qquad (13)$$

where $N(v,p)$ denotes the nodes that are in the *p-hop neighborhood* of v. That is, $u \in N(v,p)$ is at most p hops away from v.

Alternatively, we also provide the *global stress* function as the objective for each agent and rewrite Equation (13) as

$$E^v(t) = \sum_{u < v} w_{u,v}(d_t(u,v) - d'_t(u,v))^2. \qquad (14)$$

The reward function for agent v at time t is $\mathcal{R}^v = E^v(t-1) - E^v(t)$, where $E^v(t)$ is either the local or the global stress function.

4.3 Derive New Algorithms with Aesthetic Criteria

A main advantage of our MARL model for graph drawing is its generality, that is, the ability to create new layout algorithms by customizing, for each agent, its reward function, as well as its state and action spaces. In the most general form, a reward function $\mathcal{R}^v$ on an agent v at time t can be used not only to encode forces and energy on the node (as shown in Sect. 4.2), but also to incorporate various aesthetic criteria for graph layouts.

On a highlevel, a quality measure Q for a graph layout is defined to be a linear combination of aesthetic criteria,

$$Q = \sum_i \omega_i q_i, \quad \sum_i \omega_i = 1 \qquad (15)$$

where q_i is the i-th aesthetic criterion, ω_i is its coefficient. Potential aesthetic criteria include the number of edge crossings, the amount of node overlaps, the distance between non-adjacent nodes, edge length variance, and the cumulative deviation of edge angles from the ideal values. These aesthetic criteria have been widely used in the evolutionary graph layout algorithms (e.g., [3, 32, 41, 49]). The coefficients could be determined empirically [49] or learned using

genetic programming [32], as Q is related (but not equivalent) to the fitness of a graph layout.

In this paper, we experiment with a few aesthetic criteria from [49] and [32] below. The main idea is to define a local quality measure Q^v to be associated with each agent v, and subsequently to encode Q^v in the reward function $\mathcal{R}^v$.

Edge length variance. For a given graph layout, the *edge length variance* measures the mean relative square error of edge lengths [49], $\sigma = \frac{1}{|E|} \sum_{e \in E} \left(\frac{\|e\| - L}{L} \right)^2$, where $\|e\|$ is the length of an edge and L is the desired edge length. σ can be modified locally for a single agent v, $\sigma^v = \frac{1}{|E_v|} \sum_{e \in E_v} \left(\frac{\|e\| - L}{L} \right)^2$, where E_v represents edges incident to node v.

Edge angle deviation. The *edge angle deviation* captures the cumulative square deviation Δ of edge angles from their ideal values [49], $\Delta = \sum_{v \in V} \sum_{k=1}^{|E_v|} \left(\psi_k(v) - \frac{2\pi}{|E_v|} \right)^2$, where E_v contains edges incident to v, and $\psi_k(v)$ are the angles between adjacent edges in E^v. For an agent v, we can derive its local edge angle deviation as $\Delta^v = \sum_{k=1}^{|E_v|} \left(\psi_k(v) - \frac{2\pi}{|E_v|} \right)^2$.

Node overlaps. For a given graph layout, let φ represents the total number of node overlaps. Let $u < v$ denote an ordering among distinct vertices $u, v \in V$.

$$\varphi = \sum_{u \in V, v \in V, u < v} \begin{cases} 1, & \text{if } u \text{ and } v \text{ overlap;} \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

For an agent v, the number of node overlaps local to v, denoted as φ^v, is the number of nodes in V that overlaps with v,

$$\varphi^v = \sum_{u \in V, u \neq v} \begin{cases} 1, & \text{if } u \text{ overlaps with } v; \\ 0, & \text{otherwise.} \end{cases} \quad (17)$$

Edge crossings. Let χ represent the total number of edge crossings. For agent v, the number of edge crossings local to v is defined as

$$\chi^v = \sum_{e_i \in E_v} \sum_{e_j \in E} \begin{cases} 1, & \text{if } e_i \text{ and } e_j \text{ cross in their interior;} \\ 0, & \text{otherwise;} \end{cases} \quad (18)$$

where E_v represents edges incident to node v.

Minimum distance between nodes. The aesthetic criteria η ensures that there is enough space between nodes [32]. For agent v, the local version of η is defined as

$$\eta^v = \sum_{u \in V} \begin{cases} L - d(u, v), & \text{if } d(u, v) < L; \\ 0, & \text{otherwise;} \end{cases} \quad (19)$$

where L is the expected minimum distance between nodes.

Local quality measure. Given the above aesthetic criteria, we derive a local quality measure for an agent v as

$$Q^v = \omega_1 \varphi^v + \omega_2 \chi^v + \omega_3 \eta^v + \omega_4 \sigma^v + \omega_5 \Delta^v. \quad (20)$$

We then define a reward function $\mathcal{R}^v$,

$$\mathcal{R}^v(t) = Q^v(t-1) - Q^v(t), \quad (21)$$

based upon Q^v evaluated at time $t-1$ and t. Notice that Equation (21) takes a similar form as Equation (9) and Equation (10).

4.4 Extensions to Incremental and Hybrid Layouts

Finally, our MARL model can be employed to interpret incremental graph layout algorithms. In incremental layouts, instead of computing a new layout for the entire graph, only a small part of the graph is rearranged, and the positions of the remaining nodes stay (roughly) the same [34, 35]. Incremental layouts are important for maintaining a mental map of the users, which helps them keep track of changes in the graph [37]. Incremental layouts can also be used for online dynamic graphs [2, 9].

Our MARL model can be modified in various ways to act as an incremental layout algorithm. One obvious approach would be to lock a node (or nodes) that we do not want to move by limiting the action space of its agent (to only contain the "stay" action). Another approach is to modify the reward function, that is, nodes are still allowed to move, but any action other than the "stay" action has a lower reward. Restricting the state and action spaces of our MARL model will allow us to create an incremental version of most classical graph drawing algorithms; see the supplementary video for a demo.

Finally, an additional advantage of our MARL model is that the reward function is highly customizable. For instance, a hybrid reward function can be created by taking a linear combination of force-based, energy-based, and/or aesthetic-based rewards, which is difficult to do with classic layout algorithms. As an example, we can create a new layout algorithm that is a hybrid of a FR layout and stress majorization:

$$\mathcal{R}^v(t) = \beta(F^v(t-1) - F^v(t)) + (1 - \beta)(E^v(t-1) - E^v(t)), \quad (22)$$

where β is the weight, $F^v(t)$ and $E^v(t)$ are given in Equation (11) and Equation (13), respectively.

5 ALGORITHM AND IMPLEMENTATION

The proposed MARL framework for graph drawing is sketched in algorithm 1. This general framework is applicable to the MARL versions for DCG layout, FR layout, and stress majorization as well as to MARL layouts with hybrid and customized rewards.

Algorithm 1: MARL Graph Layout

   input : A graph $G(E, V)$, a reward function $\mathcal{R}$
   output : The position of each node $v \in V$

1 // initialization
2 **for** $v \in V$ **do**
3   create and assign an agent to node v;
4   initialize the location of v randomly;

5 // iterative layout procedure
6 **while** *not converged* **do**
7   **for** $v \in V$ **do**
8   // obtain reward for the agent
9         $\mathcal{R}^v \leftarrow$ calculateReward(v);
10   // obtain the next action for the agent
11        $a^v \leftarrow$ getNextAction(v);
12   // update the location of node
13        $v \leftarrow$ moveNode(v, a^v);
14   // update the reward of the agent
15        $\mathcal{R}^v \leftarrow \mathcal{R}^v$ - calculateReward(v);
16   updateRewardTable(v, $\mathcal{R}^v$);

The calculateReward function evaluates the force-based, energy-based, or customized reward for a given agent. The getNextAction function returns the next action based on the greedy policy of an agent. Once the next action is obtained, the node controlled by an agent is moved in the direction of the given action via

moveNode. The updateRewardTable updates the Q-function values based on Equation (8).

The proposed MARL framework API is implemented as a Cytoscape.js [14] extension in JavaScript in combination with REINFORCEjs[2], a reinforcement learning library. For reproducibility, parameter settings for all graph drawing algorithms described in this paper is detailed in Appendix A.

6 A VISUAL DEMO VIA MARLL

In addition to an API, we also provide an open-source, interactive tool called MARLL. MARLL demonstrates the capabilities of our MARL algorithm, evaluates and compares the resulting graph layouts against classic layout algorithms. It is built with HTML/CSS/Javascript stack with Cytoscape.js, D3.js and JQuery Javascript libraries. MARLL allows users to experiment with the MARL framework by providing their own graph files or choosing among existing sample graphs. Its user interface is shown in Fig. 2.
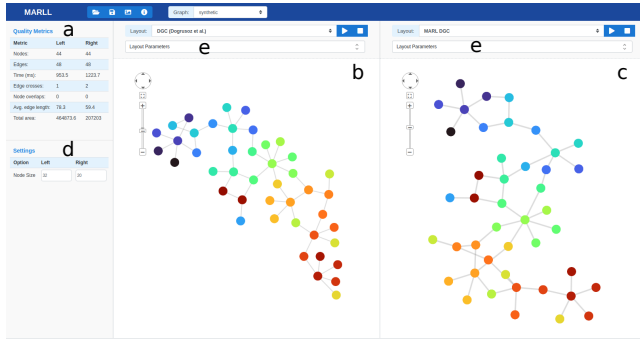

Figure 2: The interactive user interface of MARLL.

The user interface consists of three main panels. Panel (a) contains statistics and quality metrics used to evaluate current graph layouts quantitatively. To facilitate layout comparisons, the interface provides two adjacent linked views – panel (b) and panel (c) – that visualize a graph with two different layout algorithms side-by-side. Elements (nodes and edges) highlighted in panel (b) will be highlighted in panel (c), and vice versa. Panel (d) allows users to change the sizes of nodes in either view. Users can also use panel (e) to modify the parameters of the selected layout algorithms. The nodes are colored based on their indices using Turbo, an improved rainbow [33] colormap, which helps identify corresponding nodes between the two views.

7 EXPERIMENTS AND EVALUATION

In this section, we perform quantitative and qualitative evaluations together with runtime analysis. We demonstrate via experiments that MARL graph layout algorithms produce results that are aesthetically comparable to those of the classical layout algorithms. At the same time, MARL layout algorithms are easily generalizable and provide a unifying framework for a set of classic layout algorithms.

7.1 Datasets

For our experiments, we use 10 graphs with varying types and sizes. The graphs, descriptions, and sources are given in Table 1.

7.2 Parameter Configurations for Evaluation Purposes

We implement and evaluate classic graph layout algorithms in comparison with their MARL counterparts.

Specifically, we include force-directed layouts such as **FR** layout by Fruchterman and Reingold [15] and **DGC** layout by Dogrusoz et al. [12]; their MARL versions are denoted as **MARL FR** and

Table 1: The graphs used in our experiments.

Graph	Name	Type	$|V|$	$|E|$	Source
G_1	karate	social network	34	78	[29, 54]
G_2	bcspwr01	power system	39	46	[11]
G_3	synthetic	synthetic	44	48	[12]
G_4	lesmiserables	social network	77	254	[24, 29]
G_5	GD99_c	graph drawing contest	105	149	[11]
G_6	bcspwr03	power system	118	179	[11]
G_7	rajat05	circuit simulation	301	952	[11]
G_8	GD00_c	graph drawing contest	638	1031	[11]
G_9	crime_moreno	social network	829	1473	[11]
G_{10}	maayan-faa	flight network	1226	2613	[42]

MARL DGC, respectively. The **MARL Custom** layout is based on local quality measure in Equation (21). For a reasonable comparison, these layout algorithms are ran until convergence with the same convergence criteria, that is, with displacement rate $\Delta A = 2$ (pixels). We also include stress majorization (**SM**) by Gansner et al. [17], whose MARL counterparts are denoted as **MARL Global Stress** and **MARL Local Stress**, respectively, based on Sect. 4. Their shared convergence criterion is the stress ratio $\delta E = 0.0001$. Finally, we create a hybrid reward function in the MARL setting (**MARL Hybrid**), where the algorithm converges based on both displacement rate and stress ratio (whichever comes first).

For each of the layout algorithm, the node positions are initialized randomly. We run each layout algorithm 100 times and report the average evaluation measures (such as aesthetic criteria and runtime). The evaluation is conducted on a laptop machine with an Intel i7-9850H processor (6 cores at 2.6 GHz) CPU and 64 GB of RAM.

7.3 Quantitative Evaluation

To evaluate the layouts quantitatively, we employ a list of aesthetic criteria from the literature (e.g. [36, 38, 49]), namely, minimizing the number of node overlaps, reducing the number of edge crossings [36], maximizing the minimum angle between edges leaving a node [36] (similar to the the *edge angle deviation* [49]), and minimizing the mean relative square error of edge lengths [49] (the *edge length variance*); see Sect. 4.3 for their definitions.

For comparative purposes, the measurement for each aesthetic criteria is normalized to be a real number in $[0,1]$ following the metrics proposed by Purchase [36], where a higher value indicates a better layout aesthetically. Such a normalization ensures that "the metric value does not depend on the nature of the underlying graph" [36].

Normalized edge crossings (NC). Recall $|V| = n$ and $|E| = m$. Let χ denote the number of edge crossings in a layout. Let χ^* denote an approximation for the upper bound on the number of edge crossings [36],

$$\chi^* = \frac{m(m-1)}{2} - \frac{1}{2}\sum_{i=1}^{n} \deg(v_i)(\deg(v_i) - 1). \quad (23)$$

In Equation (23), the first component is the maximum possible edge crossings, and the second component is the total number of impossible crossings (since adjacent edges cannot cross).

The normalized number of edge crossings (abbreviated as **NC**) is based on the *edge crossings aesthetic metric* of Purchase [36],

$$N_c = 1 - \begin{cases} \frac{\chi}{\chi^*}, & \text{if } \chi^* > 0; \\ 0, & \text{otherwise.} \end{cases} \quad (24)$$

Normalized node overlaps (NO). Let $\binom{n}{2} = n(n-1)/2$ denote the maximum number of node overlaps in a layout. Recall φ is the

total number of node overlaps, see Equation (16). We define the *normalized number of node overlaps* (abbreviated as **NO**) as

$$N_o = 1 - \frac{\varphi}{n(n-1)/2}. \tag{25}$$

Normalized edge lengths (NE). The normalized edge lengths (abbreviated as **NE**) measure how uniform the edge lengths are. We define NE as

$$N_e = \frac{1}{1+\sigma}$$

where σ is the mean relative square error of edge lengths (the edge length variance).

Normalized edge angles (NA). The normalized edge angle (abbreviated as **NA**) is based on the *minimum angle aesthetic metric* [36]. It measures how well the edge angles are distributed surrounding a given node. For instance, for a degree-4 node, the ideal angle is 90 degree. The normalized edge angle is defined as

$$N_a = 1 - \frac{1}{n} \sum_{i=1}^{n} \left| \frac{\theta_i^* - \theta_i}{\theta_i^*} \right|, \tag{26}$$

where θ_i^* is the ideal (maximal) minimum angle at the *i*-th node, $\theta_i^* = 360^\circ / \deg(v_i)$; and θ_i is the actual minimum angle between the incident edges at the *i*-th node [36].

Results. We use these four metrics (**NC, NO, NE,** and **NA**) to perform quantitative evaluations. The results are given in Table 2 and Table 3, respectively.

In Table 2 , R_{NC}, R_{NO}, R_{NE}, and R_{NA} represent the ratio of measurements for the MARL version to the classical version, respectively. For example, R_{NC} represents the ratio of the number of edge crossings for MARL version to the number of edge crossings for the classical version. The higher the ratios are, the more aesthetically comparable the results are between the MARL layouts and the classical layouts.

In Table 3, we give quantitative evaluation results using a local quality measure as the custom reward function (Equation (21)). The results indicate that our MARL framework with a custom reward function produces aesthetically comparable layouts in comparison with other algorithms, shown in Table 2.

The quantitative evaluations show that the MARL layout algorithms produce results very similar to those of their classical counterparts in terms of the four quality metrics, independent of the graph size. The two MARL layout algorithms based on the custom and hybrid reward functions do not have a classical version for a direct comparison. However, their quality metrics are close to those for classical layout algorithms.

7.4 Qualitative Evaluation

We also perform a qualitative evaluation to demonstrate that the MARL framework produces results aesthetically comparable to those of the classical layout algorithms. We show examples G_4, G_6, and G_7 in Fig. 3, where we compare the layouts obtained by classical algorithms (DGC, FR, and stress majorization) with those obtained by their MARL versions. We also include layouts obtained via a MARL hybrid and MARL custom. For each graph, nodes with the same color correspond to one another. Due to space constraints, additional layouts for all graphs can be found in Appendix B.

Based on our qualitative evaluation, the classical layout algorithms and the MARL versions produce graph drawings that are visually alike. For some of the graphs, the results are visually indistinguishable. The layouts of the MARL hybrid layout are as good as the layouts from the local and global stress MARL algorithm. Since the MARL layout with a custom reward function does not have a classical counterpart, its outputs are not expected to be visually similar to others.

7.5 Runtime Analysis

Finally, we compare the runtime of classic layout algorithms until convergence against their corresponding MARL versions; see Table 4 for force-directed layout algorithms and Table 5 for stress majorization. We follow the parameter configurations detailed in Sect. 7.2. When applicable, we report the ratio of runtime of MARL layouts to their classical counterparts under the same convergence criteria. All runtimes are reported in milliseconds.

The runtimes for the MARL layouts with custom (Equation (21)) and hybrid (Equation (22)) reward functions are provided in Table 6.

In comparison to the classical algorithms, the runtime overhead for MARL versions is very large for smaller sized graphs, especially for the MARL layout based on [12] and [17]. However, as the graph size increases, the difference between the run-time of classical and their MARL versions decreases. The MARL Hybrid has a computation time similar to that of the MARL layouts based on local stress and global stress. Although the MARL Custom does not have a classical version for comparison, its generally has the slowest computation time among all MARL layouts.

8 DISCUSSION

In this paper, we propose a novel MARL-based approach to graph drawing. Our approach interprets several classical force-directed and energy-based graph drawing algorithms with MARL. Additionally, it enables derivation of new layout algorithms by providing custom and hybrid reward functions to the agents of the system.

The main attraction for a MARL formulation is that it is flexible, and can be adapted to existing and new graph drawing algorithms. By modeling with MARL, we provide a unifying framework for a number of classical layout algorithms, which is also easily generalizable to create custom and hybrid layout algorithms. Our evaluation results demonstrate that the MARL graph drawing algorithms generate comparable results to their classical counterparts; with some amount of computational overhead. We hope this work will inspire more research in adapting RL for graph drawing.

We end our paper with some further discussion on scalability. It is well known that most MARL algorithms have exponential computation complexity in the joint state-action space [57]. That is, the size of the state-action space is exponentially large in the number of agents [40]. Since we utilize the MARL framework in graph drawing, our framework inherits the exponential complexity of MARL [6] in terms of the number of nodes in a graph, making the graph drawing difficult to scale to really large graphs.

However, not all hope is lost for scalability. Recent work by Qu *et al.* [39, 40] identified a class of networked MARL problems where "the model exhibits a local dependence structure that allows it to be solved in a scalable manner". In particular, if a MARL model enforces local interactions, that is, agents are "associated with a graph and they interact only with nearby agents in the graph" [39], then it is possible to utilize the graph structure to develop scalable MARL algorithms [39, 40]. This *localized policy* fits well with our MARL graph drawing framework. In this paper, we focus on showing the feasibility of MARL for graph drawing; we leave it for future work to address its scalability challenge.

REFERENCES

[1] H. J. C. Barbosa and A. M. S. Barreto. An interactive genetic algorithm with co-evolution of weights for multiobjective problems. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation*, pages 203–210, 2001.

[2] F. Beck, M. Burch, S. Diehl, and D. Weiskopf. A taxonomy and survey of dynamic graph visualization. *Computer Graphics Forum*, 36(1):133–159, 2017.

[3] J. Branke, F. Bucher, and H. Schmeck. Using genetic algorithms for drawing undirected graphs. *Proceedings of the 3rd Nordic Workshop on Genetic Algorithms and Their Applications*, pages 193–206, 1996.
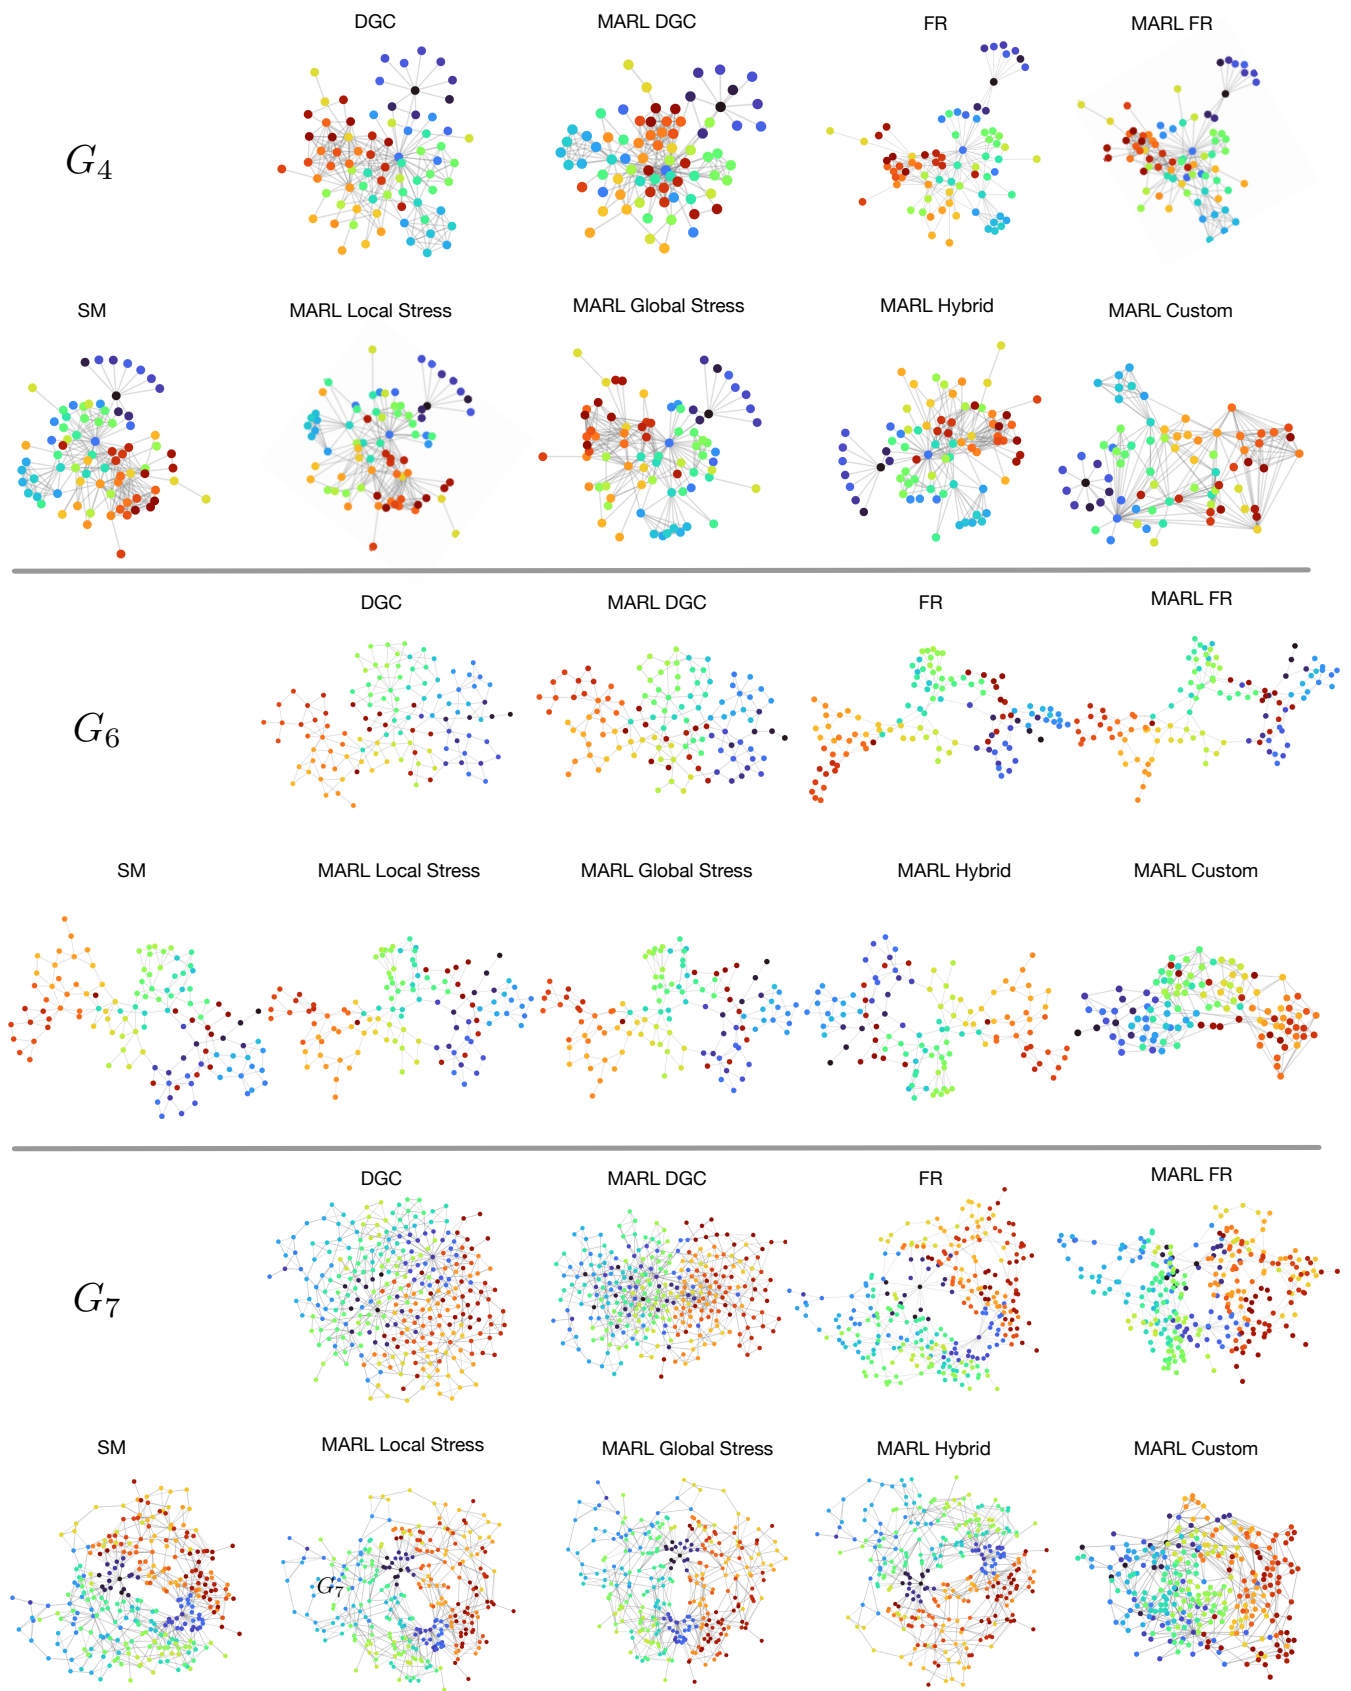
Figure 3: Qualitative results for Q_4 (top), G_6 (middle) and Q_7 (bottom) for DGC layout, FR layout, stress majorization (SM), and their MARL versions; as well as MARL hybrid and MARL custom versions.

Table 2: Quantitative evaluation results for classical force-directed layouts, stress majorization, and their MARL versions.

Graph	DGC				MARL DGC								FR				MARL FR							
	NC	NO	NE	NA	NC	NO	NE	NA	R_{NC}	R_{NO}	R_{NE}	R_{NA}	NC	NO	NE	NA	NC	NO	NE	NA	R_{NC}	R_{NO}	R_{NE}	R_{NA}
G_1	0.96	1	0.91	0.24	0.91	0.99	0.85	0.22	**0.95**	**0.99**	**0.93**	**0.91**	0.96	1	0.93	0.25	0.95	0.97	0.86	0.23	**0.99**	**0.97**	**0.93**	**0.91**
G_2	1	1	1	0.73	0.95	1	0.93	0.65	**0.95**	**1**	**0.93**	**0.89**	1	1	1	0.73	0.99	0.99	0.91	0.64	**0.99**	**0.99**	**0.91**	**0.88**
G_3	1	1	1	0.83	0.96	1	0.96	0.77	**0.96**	**1**	**0.96**	**0.93**	1	1	1	0.81	0.99	0.99	0.9	0.73	**0.99**	**0.99**	**0.9**	**0.9**
G_4	0.96	1	0.86	0.35	0.92	0.99	0.8	0.33	**0.96**	**0.99**	**0.92**	**0.93**	0.97	1	0.89	0.37	0.95	0.98	0.79	0.34	**0.99**	**0.98**	**0.89**	**0.92**
G_5	1	1	0.99	0.72	0.96	1	0.92	0.64	**0.96**	**1**	**0.92**	**0.88**	1	1	0.99	0.71	0.98	0.99	0.89	0.62	**0.98**	**0.99**	**0.9**	**0.87**
G_6	1	1	0.98	0.54	0.95	0.99	0.91	0.5	**0.95**	**0.99**	**0.94**	**0.93**	1	1	0.97	0.54	0.97	0.99	0.87	0.48	**0.98**	**0.99**	**0.89**	**0.89**
G_7	0.99	1	0.9	0.34	0.93	0.99	0.82	0.32	**0.93**	**0.99**	**0.91**	**0.93**	0.99	1	0.9	0.34	0.97	0.99	0.84	0.31	**0.98**	**0.99**	**0.92**	**0.93**
G_8	0.99	1	0.88	0.61	0.89	0.99	0.81	0.55	**0.9**	**0.99**	**0.92**	**0.89**	0.99	1	0.88	0.62	0.94	0.99	0.86	0.56	**0.95**	**0.99**	**0.98**	**0.9**
G_9	0.97	1	0.88	0.5	0.87	0.99	0.81	0.44	**0.9**	**1**	**0.92**	**0.88**	0.97	1	0.89	0.5	0.93	0.99	0.87	0.45	**0.95**	**0.99**	**0.98**	**0.9**
G_{10}	0.99	1	0.87	0.43	0.85	1	0.84	0.37	**0.86**	**1**	**0.97**	**0.85**	0.99	1	0.87	0.44	0.92	0.99	0.84	0.39	**0.93**	**0.99**	**0.96**	**0.89**

Graph	SM				MARL Local Stress								MARL Global Stress								MARL Hybrid			
	NC	NO	NE	NA	NC	NO	NE	NA	R_{NC}	R_{NO}	R_{NE}	R_{NA}	NC	NO	NE	NA	R_{NC}	R_{NO}	R_{NE}	R_{NA}	NC	NO	NE	NA
G_1	0.96	1	0.93	0.25	0.95	0.99	0.86	0.22	**0.99**	**0.99**	**0.92**	**0.88**	0.95	0.99	0.86	0.22	**0.99**	**0.99**	**0.92**	**0.87**	0.95	0.99	0.86	0.22
G_2	1	1	1	0.73	0.99	1	0.91	0.67	**0.99**	**1**	**0.91**	**0.92**	0.99	1	0.9	0.66	**0.99**	**1**	**0.91**	**0.91**	0.99	1	0.91	0.67
G_3	1	1	1	0.81	0.99	1	0.92	0.77	**0.99**	**1**	**0.92**	**0.94**	0.99	1	0.92	0.76	**0.99**	**1**	**0.93**	**0.94**	0.99	1	0.92	0.76
G_4	0.97	1	0.89	0.37	0.95	0.99	0.83	0.32	**0.98**	**0.99**	**0.93**	**0.87**	0.95	0.99	0.83	0.32	**0.98**	**0.99**	**0.93**	**0.88**	0.94	0.99	0.84	0.32
G_5	1	1	0.99	0.71	0.99	1	0.89	0.65	**0.99**	**1**	**0.9**	**0.92**	0.99	1	0.91	0.65	**0.99**	**1**	**0.92**	**0.91**	0.99	1	0.9	0.65
G_6	1	1	0.97	0.54	0.98	1	0.9	0.49	**0.99**	**1**	**0.93**	**0.91**	0.99	1	0.88	0.48	**0.99**	**1**	**0.9**	**0.9**	0.98	1	0.9	0.5
G_7	0.99	1	0.9	0.34	0.99	1	0.81	0.3	**1**	**1**	**0.9**	**0.88**	0.99	1	0.82	0.29	**0.99**	**1**	**0.9**	**0.87**	0.99	1	0.82	0.3
G_8	0.99	1	0.88	0.62	0.98	1	0.79	0.58	**0.99**	**1**	**0.89**	**0.93**	0.98	1	0.79	0.57	**0.99**	**1**	**0.89**	**0.93**	0.98	1	0.79	0.58
G_9	0.97	1	0.89	0.5	0.96	1	0.8	0.46	**0.99**	**1**	**0.9**	**0.91**	0.96	1	0.81	0.46	**0.99**	**1**	**0.92**	**0.91**	0.96	1	0.78	0.46
G_1	0.99	1	0.87	0.44	0.98	1	0.76	0.37	**0.99**	**1**	**0.87**	**0.86**	0.98	1	0.77	0.37	**0.99**	**1**	**0.88**	**0.86**	0.98	1	0.77	0.37

Table 3: Quantitative evaluation results for MARL layouts based on custom reward function.

Graph	MARL Custom Reward			
	NC	NO	NE	NA
G_1	0.931	0.989	0.875	0.243
G_2	0.963	0.992	0.896	0.581
G_3	0.96	0.993	0.859	0.687
G_4	0.931	0.993	0.779	0.357
G_5	0.951	0.994	0.849	0.591
G_6	0.936	0.993	0.825	0.456
G_7	0.962	0.996	0.778	0.29
G_8	0.91	0.996	0.767	0.559
G_9	0.913	0.996	0.753	0.459
G_{10}	0.937	0.997	0.774	0.397

Table 4: Runtimes in milliseconds for classical and MARL force-directed layouts.

Graph	Runtime of Force Directed Layouts					
	DGC	MARL DGC	Ratio DGC	FR	MARL FR	Ratio FR
G_1	72	438	6.05	140	352	2.51
G_2	74	458	6.17	157	397	2.53
G_3	85	582	6.85	304	481	1.58
G_4	285	1744	6.11	606	1011	1.67
G_5	276	1953	7.08	705	886	1.26
G_6	328	2563	7.81	884	1008	1.14
G_7	3598	16810	4.67	4030	5052	1.25
G_8	15383	68801	4.47	13688	18597	1.36
G_9	34347	82755	2.41	26890	32378	1.2
G_{10}	70100	143892	2.05	54684	68022	1.24

Table 5: Runtimes in milliseconds for classical stress majorization algorithm and its MARL variants.

Graph	Runtime of Stress Majorization Layouts				
	SM	MARL Local Stress	Ratio Local Stress	MARL Global Stress	Radio Global Stress
G_1	17	126	7.46	135	7.99
G_2	20	177	8.88	187	9.36
G_3	25	249	9.95	257	10.27
G_4	64	398	6.42	386	6.23
G_5	50	550	10.96	605	12.07
G_6	65	630	9.74	633	9.79
G_7	544	5154	9.48	5235	9.63
G_8	6700	25696	3.84	26050	3.89
G_9	10884	42274	3.88	42730	3.93
G_{10}	36360	118486	3.26	121637	3.35

Table 6: Runtimes in milliseconds for the MARL layouts with custom and hybrid reward functions.

Graph	Runtime of Hybrid and Custom	
	MARL Hybrid	MARL Custom
G_1	128	394
G_2	186	454
G_3	269	474
G_4	430	1552
G_5	702	1110
G_6	746	1725
G_7	6304	12112
G_8	30764	126504
G_8	45681	232724
G_{10}	138034	752484

[4] N. Brown and T. Sandholm. Libratus: The superhuman ai for no-limit Poker. *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 5226–5228, 2017.

[5] L. Busoniu, R. Babuska, and B. De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.

[6] L. Buşoniu, R. Babuška, and B. D. Schutter. Multi-agent reinforcement learning: An overview. In D. Srinivasan and L. C. Jain, editors, *Innovations in Multi-Agent Systems and Applications - 1. Studies in Computational Intelligence*, pages 183–221. Springer, Berlin, Heidelberg, 2010.

[7] R. Case. Machine learning, part 4: Reinforcement learning. https://towardsdatascience.com/, 2019.

[8] A. Cimikowski and P. Shope. A neural-network algorithm for a graph layout problem. *IEEE Transactions on Neural Networks*, 7(2):341–345, 1996.

[9] T. Crnovrsanin, J. Chu, and K.-L. Ma. An incremental layout method for visualizing online dynamic graphs. In E. Di Giacomo and A. Lubiw, editors, *Graph Drawing and Network Visualization*, Lecture Notes in Computer Science, pages 16–29, Cham, 2015. Springer International Publishing.

[10] R. Davidson and D. Harel. Drawing graphs nicely using simulated annealing, Oct. 1996.

[11] T. A. Davis and Y. Hu. The university of Florida sparse matrix collection, Dec. 2011.

[12] U. Dogrusoz, E. Giral, A. Cetintas, A. Civril, and E. Demir. A layout algorithm for undirected compound graphs. *Information Sciences*, 179(7):980–994, 2009.

[13] P. Eades. A heuristic for graph drawing. *Congressus Numerantium*, 42:149–160, 1984.

[14] M. Franz, C. T. Lopes, G. Huck, Y. Dong, O. Sumer, and G. D. Bader. Cytoscape.js: A graph theory library for visualisation and analysis. *Bioinformatics*, 32(2):309–311, 2016.

[15] T. M. J. Fruchterman and E. M. Reingold. Graph drawing by force-directed placement. *Software—Practice & Experience*, 21(11):1129–1164, 1991.

[16] E. R. Gansner. Using Graphviz as a library (cgraph version). https://www.graphviz.org/pdf/libguide.pdf, 2014.

[17] E. R. Gansner, Y. Koren, and S. North. Graph drawing by stress majorization. In J. Pach, editor, *Graph Drawing*, pages 239–250. Springer, 2005.

[18] H. Gibson, J. Faith, and P. Vickers. A survey of two-dimensional graph layout techniques for information visualisation. *Information Visualization*, 12(3-4):324–357, 2013.

[19] I. Herman, G. Melancon, and M. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, Jan. 2000.

[20] Y. Hu. Efficient, high-quality force-directed graph drawing. *The Mathematica Journal*, page 35, 2006.

[21] L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[22] T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Information Processing Letters*, 31(1):7–15, 1989.

[23] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *Proceedings of International Conference on Learning Representations*, 2014.

[24] D. E. Knuth. *The Stanford GraphBase: A Platform for Combinatorial Computing*. Association for Computing Machinery, New York, NY, USA, 1993.

[25] J. Kober and J. Peters. *Learning Motor Skills: From Algorithms to Robot Experiments (Springer Tracts in Advanced Robotics (97))*. Springer International Publishing, 2014.

[26] S. G. Kobourov. Spring embedders and force directed graph drawing algorithms. arXiv:1201.3011, 2012.

[27] C. Kosak, J. Marks, and S. Shieber. Automating the layout of network diagrams with specified visual organization. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(3):440–454, 1994.

[28] C. Kosak, J. Marks, and S. M. Shieber. A parallel genetic algorithm for network-diagram layout. *Proceedings of the 4th International Conference on Genetic Algorithms*, 1991.

[29] J. Kunegis. KONECT: The Koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, WWW '13 Companion, pages 1343–1350, Rio de Janeiro, Brazil, May 2013. Association for Computing Machinery.

[30] O.-H. Kwon, T. Crnovrsanin, and K.-L. Ma. What would a graph look like in this layout? a machine learning approach to large graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):478–488, 2018.

[31] O.-H. Kwon and K.-L. Ma. A Deep Generative Model for Graph Layout. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):665–675, Jan. 2020.

[32] T. Masui. Evolutionary learning of graph layout constraints from examples. *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology*, pages 103–108, 1994.

[33] A. Mikhailov. Turbo, an improved rainbow colormap for visualization. Google AI Blog, 2019.

[34] K. Miriyala, S. Hornick, and R. Tamassia. An incremental approach to aesthetic graph layout. In *Proceedings of 6th International Workshop on Computer-Aided Software Engineering*, pages 297–308, 1993.

[35] S. C. North. Incremental layout in DynaDAG. In F. J. Brandenburg, editor, *Graph Drawing*, Lecture Notes in Computer Science, pages 409–418, Berlin, Heidelberg, 1996. Springer.

[36] H. C. Purchase. Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing*, 13(5):501–516, 2002.

[37] H. C. Purchase, E. Hoggan, and C. Görg. How important is the "mental map"? – an empirical investigation of a dynamic graph layout algorithm. In M. Kaufmann and D. Wagner, editors, *Graph Drawing*, Lecture Notes in Computer Science, pages 184–195, Berlin, Heidelberg, 2007. Springer.

[38] H. C. Purchase, C. Pilcher, and B. Plimmer. Graph drawing aesthetics – created by users, not algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):81–92, 2012.

[39] G. Qu, Y. Lin, and A. Wierman. Scalable multi-agent reinforcement learning for networked systems with average reward. *arXiv:2006.06626*, 2020.

[40] G. Qu, A. Wierman, and N. Li. Scalable reinforcement learning of localized policies for multi-agent networked systems. *Proceedings of Machine Learning Research, 2nd Annual Conference on Learning for Dynamics and Control*, 120:1–11, 2020.

[41] A. Rosete-Suarez, M. Sebag, and A. Ochoa-Rodriguez. A study of evolutionary graph drawing. Technical report, Laboratoire de Recherche en Informatique (LRI), Universite Paris-Sud, 1999.

[42] R. A. Rossi and N. K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015.

[43] S. Shalev-Shwartz, S. Shammah, and A. Shashua. Safe, multi-agent, reinforcement learning forautonomous driving. arXiv:1610.03295, 2016.

[44] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[45] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[46] M. Spönemann, B. Duderstadt, and R. von Hanxleden. Evolutionary meta layout of graphs. In T. Dwyer, H. Purchase, and A. Delaney, editors, *Diagrammatic Representation and Inference, Lecture Notes in Computer Science*, volume 8578, pages 16–30, Berlin, Heidelberg, 2014. Springer.

[47] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2012.

[48] R. M. Tarawaneh, P. Keller, and A. Ebert. A General Introduction To Graph Visualization Techniques. In C. Garth, A. Middel, and H. Hagen, editors, *Proceedings of IRTG 1131 – Visualization of Large and Unstructured Data Sets Workshop 2011*, volume 27, pages 151–164, Dagstuhl, Germany, 2012. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[49] A. G. B. Tettamanzi. Drawing graphs with evolutionary algorithms. In I. C. Parmee, editor, *Adaptive Computing in Design and Manufacture*, pages 325–337, London, 1998. Springer.

[50] W. T. Tutte. Convex Representations of Graphs. *Proceedings of the London Mathematical Society*, s3-10(1):304–320, 1960.

[51] R. Vieira, H. Nascimento, and W. Silva. The application of machine learning to problems in graph drawing: A literature review. *Proceedings of 7th International Conference on Information, Process, and Knowledge Management*, 2015.

[52] Y. Wang, Z. Jin, Q. Wang, W. Cui, T. Ma, and H. Qu. DeepDrawing: A deep learning approach to graph drawing. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):676–686, 2020.

[53] C. J. Watkins and P. Dayan. Technical note: Q-learning. *Machine Learning*, 8(3):279–292, 1992.

[54] W. W. Zachary. An Information Flow Model for Conflict and Fission in Small Groups. *Journal of Anthropological Research*, 33(4):452–473, 1977.

[55] K. Zhang, Z. Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *arXiv:1911.10635*, 2019.

[56] J. X. Zheng, S. Pawar, and D. F. M. Goodman. Graph Drawing by Stochastic Gradient Descent. *IEEE Transactions on Visualization and Computer Graphics*, 25(9):2738–2748, Sept. 2019.

[57] L. Zhou, P. Yang, C. Chen, and Y. Gao. Multiagent reinforcement learning with sparse interactions by negotiation and knowledge transfer. *IEEE Transactions on Cybernetics*, 47(5):1238–1250, 2017.

A IMPLEMENTATION DETAILS

We provide additional implementational details for reproducibility. The implementation of the FR layout and stress majorization is provided by Graphviz [16]. The DCG layout is available as part of Cytoscape.js. We implement MARL counterparts of these classical layout algorithms based on the formulation discussed in Sect. 4.2 and by modifying and integrating the library REINFORCEjs with Cytoscappe.js. In particular, we implement both local and global stress for stress majorization. We make a distinction between parameter settings for each layout algorithms in general and the convergence criteria for comparative analysis in Sect. 7.

FR, DCG, Stress Majorization and their MARL counterparts. For FR and MARL FR algorithms, the optimal edge length $k = 30$. The parameters of DCG and MARL DCG algorithms include: $\lambda = 30$, $\zeta = 5$, and $\mu = 5000$. For MARL Local Stress algorithm, $p = 10$.

MARL custom layouts. To derive new layout algorithms, we use formulation based on the local quality measure in Equation (21) and implement the customized reward based on Equation (20). The default weight for each metric is determined using a grid search that maximizes the local quality measure for a fixed small graph: $\omega_1 = 0.35$, $\omega_2 = 0.20$, $\omega_3 = 0.10$, $\omega_4 = 0.25$, and $\omega_5 = 0.10$. The values of desired edge length and expected minimum node distance in MARL Custom layout algorithm are set to 30.

MARL hybrid and incremental layouts. For our MARL hybrid implementation, the hybrid reward function is based on Equation (22) with $\beta = 0.5$. We also include a MARL-based incremental layout algorithm in MARLL.

Convergence parameters. In our graph layout framework, there are four parameters to configure for the convergence criterion; the number of maximum iterations M, average node displacement A, displacement rate ΔA, and the stress ratio δE.

Historically, Eades [13] asserted that all graphs reach to a minimal energy level after 100 iterations. Dogrusoz *et al.* [12] used 2500 as the number of maximum iterations in their implementation. We found that using the same value $M = 2500$ in our algorithms is a good balance between the maximum amount of running time and layout quality. We determined the threshold values of average node displacement and displacement rate empirically as $A = 5$ and $\Delta A = 2$ (pixels). Kamada and Kawai terminated their energy-based graph drawing when an energy threshold (referred to as a convergence precision) is reached [22]. However, they did not provide information on how the threshold is determined. Gansner *et al.* [17] suggested 0.0001 as a typical value for stress ratio tolerance. We use the same suggested value for our algorithms, that is, $\delta E = 0.0001$.

All these parameter values achieve aesthetically pleasing layouts in a reasonable amount of time. In addition, all convergence parameters can be configured both in the layout library API and from the user interface of the tool MARLL itself.

Node displacement with cooling schedule. Fruchterman and Reingold added the notion of "temperature" and "cooling" in their proposed algorithm [15]. The maximum amount of node displacement is limited by the temperature, and the temperature is decreased over time by using a cooling factor. The idea is that as the layout becomes better over time, smaller and smaller adjustments are needed, which is similar to a cooling process [15, 26]. Similar cooling schedule is used by Davidson and Harel [10], and Dogrusoz *et al.* [12].

Davidson and Harel used a geometric rule for adjusting the temperature: $T_{t+1} = \tau * T_t$, where τ is the cooling factor, and they set it to 0.75 to achieve a relatively rapid cooling [10]. Their temperature reduction schedule is based on the input size.

In our work, we utilize the temperature and cooling for all MARL layout algorithms to determine how much a node moves when an action is taken. Initially, each node can move by 10 pixels in each direction, and when the layout progresses the movement is gradually decreased. The cooling rule is the same one used by Davidson and Harel [10]. We reduce the temperature after every $|V| + |E|$ iterations.

B QUALITATIVE EVALUATION: COMPUTE RESULTS

We provide complete results of the qualitative evaluation to demonstrate that the MARL framework produces aesthetically comparable results to the classical layout algorithms. We show examples G_1 to G_5 in Fig. 4, as well as G_6 to G_{10} in Fig. 5 where we compare the layouts obtained by classical algorithms (DGC, FR, and stress majorization) with those obtained by their MARL versions. We also include layouts obtained via a MARL hybrid and MARL custom for all graphs in Fig. 6 . For each graph, nodes with the same color correspond with one another.
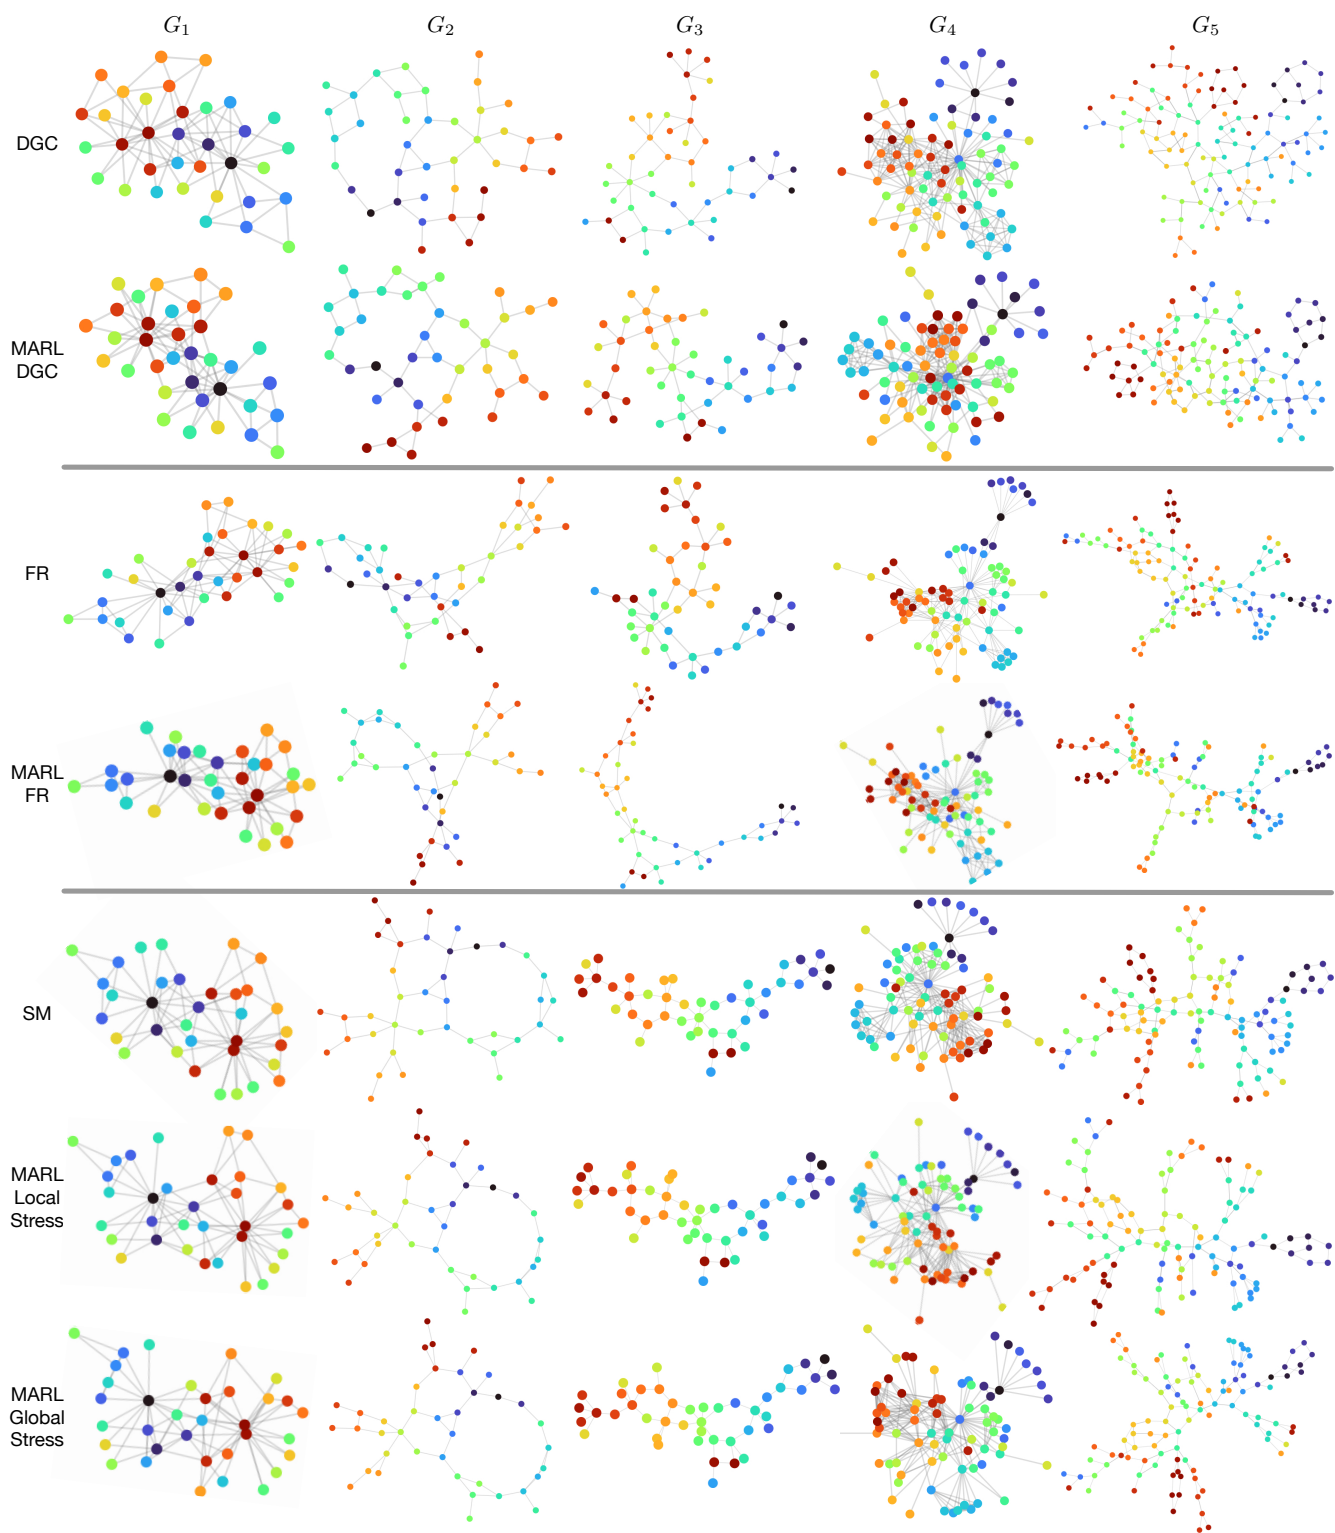
Figure 4: Qualitative results for Q_1, Q_2, Q_3, Q_4, and Q_5 for DGC layout, FR layout, stress majorization (SM), and their MARL versions.
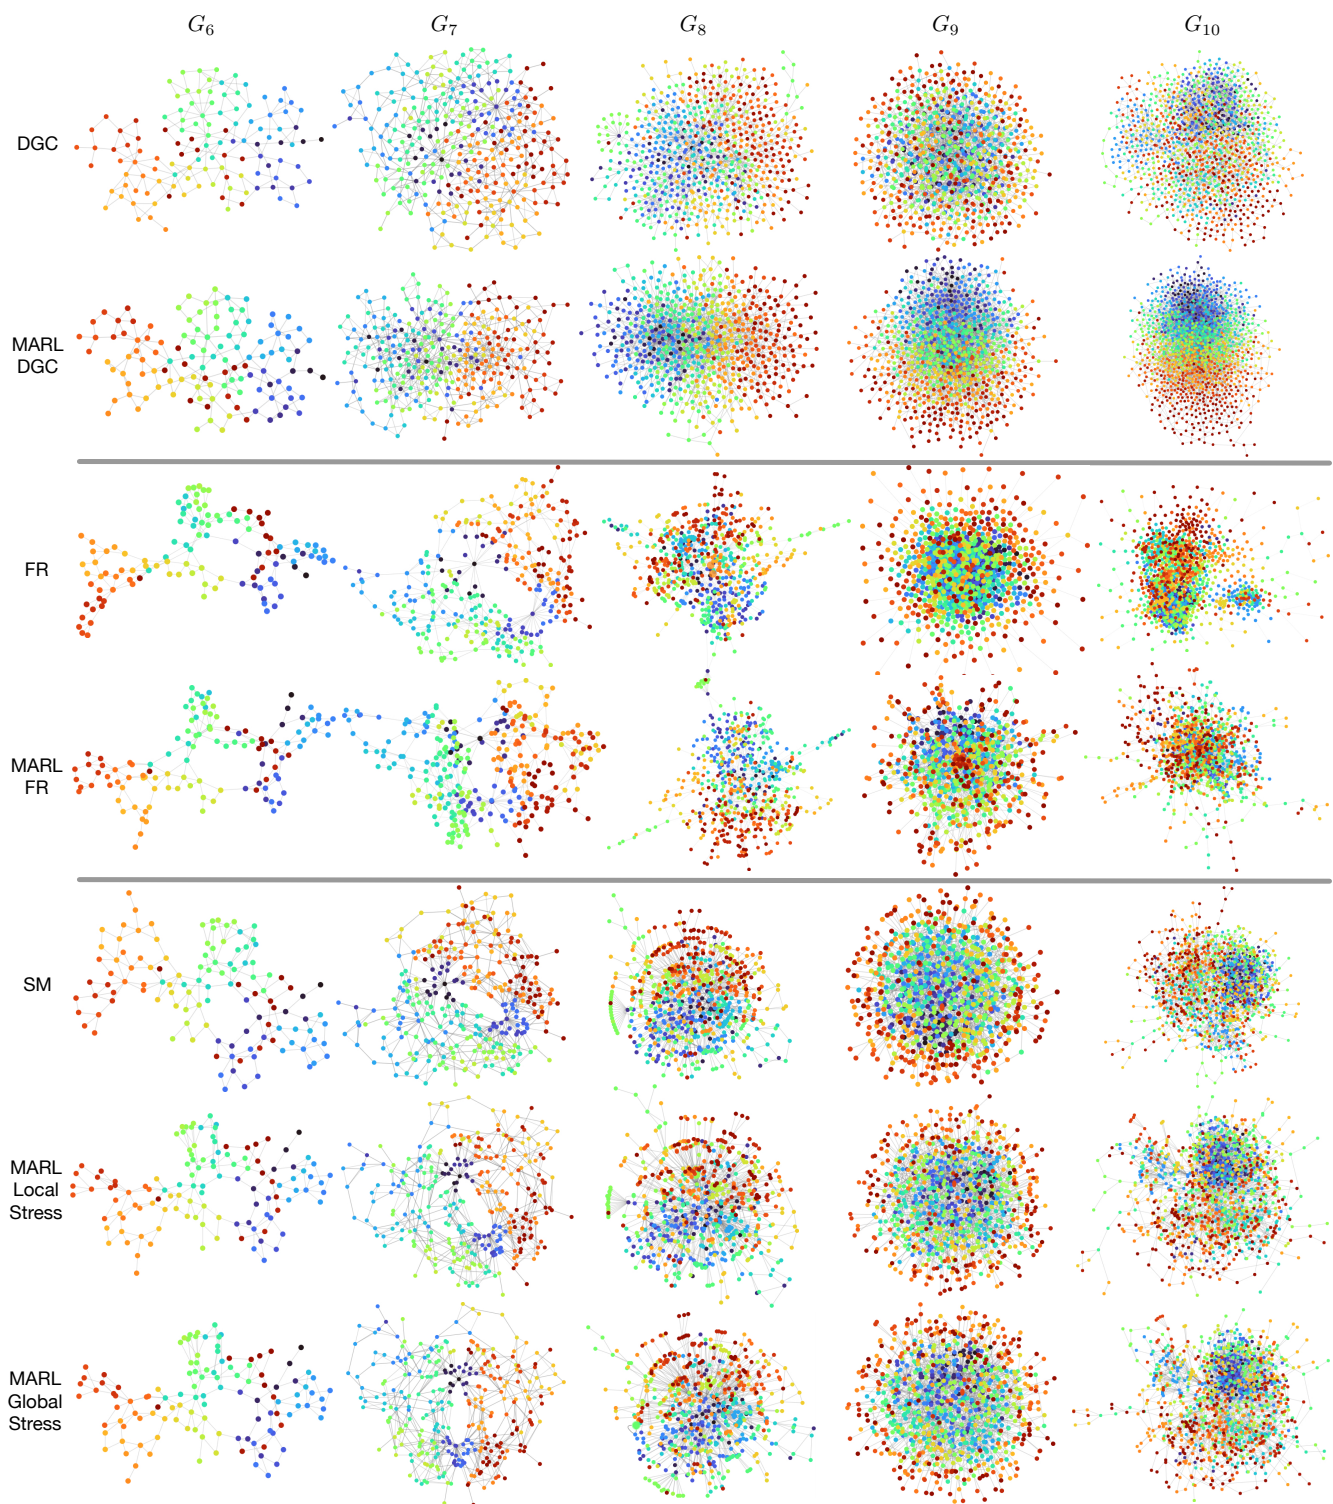
Figure 5: Qualitative results for Q_6, Q_7, Q_8, Q_9, and Q_{10} for DGC layout, FR layout, stress majorization (SM), and their MARL versions.
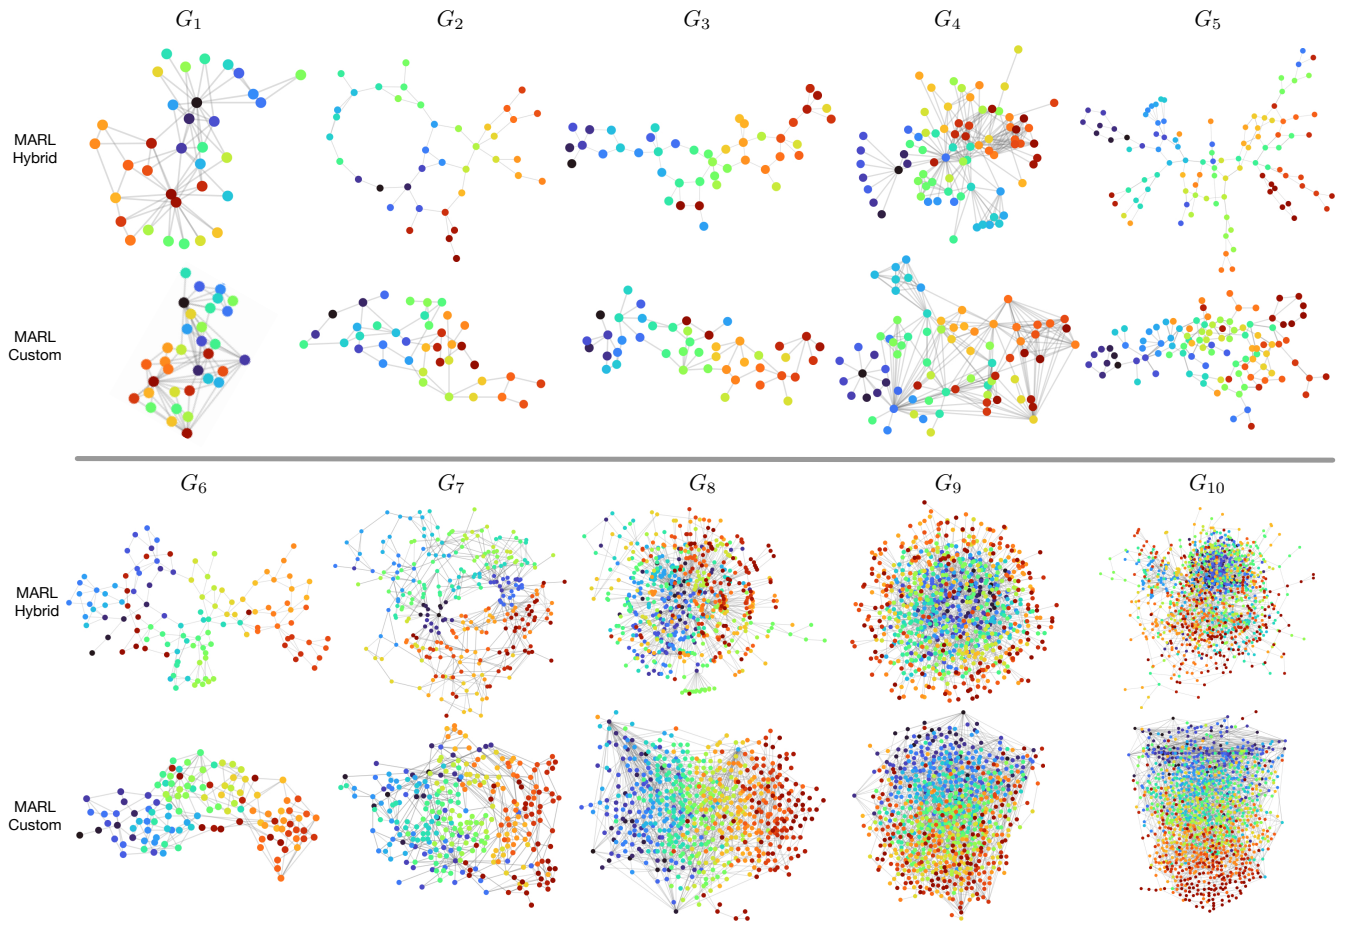
Figure 6: Qualitative results for Q_1 to Q_{10} for MARL hybrid and MARL custom versions.