



Recent Trends in Robotic Patrolling

Nicola Basilico¹

Accepted: 5 April 2022 / Published online: 6 May 2022
© The Author(s) 2022, corrected publication 2022

Abstract

Purpose of Review Robotic patrolling aims at protecting a physical environment by deploying a team of one or more autonomous mobile robots in it. A key problem in this scenario is characterizing and computing effective patrolling strategies that could guarantee some level of protection against different types of threats. This paper provides a survey of contributions that represent the recent research trends to deal with such a challenge.

Recent Findings Starting from a set of basic and recurring modeling landmarks, the formulations of robotic patrolling studied by current research are diverse and, to some extent, complementary. Some works propose optimal approaches where the objective function is based on the idleness induced by the patrolling strategy on locations of the environment. On-line methods focus on handling events that can dynamically alter the patrolling task. Adversarial methods, where an underlying game-theoretical interaction with an attacker is modeled, consider sophisticated attacker behaviors.

Summary The wide spectrum of heterogeneous approaches and techniques shows a common trend of moving towards more realistic models where constraints, dynamic environments, limited attacker capabilities, and richer strategy representations are introduced. The results provide complementarities and synergies towards more effective robotic patrolling systems, paving the way to a set of interesting open problems.

Keywords Patrolling · Surveillance · Security games · Multi-agent systems · Multi-robot systems

Introduction

The employment of mobile robots for autonomous enforcement of security represents an application domain towards which a great research effort has been devoted in the last years [1]. *Robotic patrolling* is a broad term often used to indicate those settings where one or more mobile robots persistently scout an environment visiting different parts of it and repeatedly perform local observations to identify possible ongoing threats. This general problem definition can be instantiated in different settings by specifying, for example, how the environment is represented, what robot motion and perception models are adopted, what criteria are used to evaluate the patrolling performance, and what knowledge can or must be leveraged to optimize such criteria.

The real-world implementations of such systems entail a number of engineering challenges involving the fundamental capabilities of autonomous mobile robots. These include (in increasing level of abstraction) locomotion, perception, navigation, planning, and, when multiple robots are employed, coordination [2, 3]. Achieving robust performance in these domains is a precondition to deploy the robots on the field. Patrolling shares many of these challenges with other important robotic domains, most notably exploration, coverage [4], and search [5]. Such overlaps allow for the adoption or adaptation of methods originally studied for one of these similar problems to patrolling and *vice versa*, a methodological approach that is gaining importance especially in recent research contributions.

Taking an abstract stance from the above challenges, this review will focus on a high-level, and perhaps peculiar, problem of robotic patrolling: computing effective *patrolling strategies*. A patrolling strategy is a method to determine how robots should autonomously allocate, as time unfolds, their surveillance efforts in a given environment. Roughly speaking, they repeatedly provide for each robot an answer to the question *where/how to patrol next?*, given the modeling assumptions of the patrolling setting at hand. Previous surveys on this subject have been

This article is part of the Topical Collection on *Group Robotics*

✉ Nicola Basilico
nicola.basilico@unimi.it

¹ Department of Computer Science, University of Milan, Milan, Italy

proposed in [6–8], with some examples also tailored for other research communities [9]. This work will try to complement and extend those works (also re-proposing key references under a more recent view) by sharing part of their goals and taking a robotic/multi-agent perspective.

The review will start describing a set of modeling landmarks typically used in literature to define the robotic patrolling problem (Section 2). Subsequently (Section 3) it shall provide a broad discussion about key advancements on techniques for its resolution or analysis, in the attempt of highlighting the most significant trends that recent literature is outlining. Finally, some relevant future research directions will be proposed (Section 4) before drawing conclusions (Section 5).

Problem Overview

Providing a general formulation of the robotic patrolling problem is not straightforward, especially given the remarkable number of variants that have been studied in the last two decades. Despite this, the literature suggests a set of modeling landmarks that characterize many of the problem definitions addressed so far. The frequent assumption is that the environment can be represented with a graph encoding its topology, a data structure very often (but not always) exploited for patrolling areas, lines, or perimeters (in some cases in the scope of the same method [10]). In addition, other features to describe the patrolling setting might be present. The following elements are typically adopted to model patrolling problems with graph-based representations.

Locations. $V = \{v_1, v_2, \dots, v_n\}$ is the set containing the graph vertices/nodes. Each of them is meant to represent a logical unit of space that can be visited and patrolled by a robot. Real-world counterparts depend on the specific setting and on the type of environment discretization put in place, examples can range from sub-areas and rooms of a floor plan to grid cells on a regular map [11–18, 19, 20]. One central problem is how to synthesize such a discretization from metric representations of environments. This is a problem widely studied in robotic mapping, for which robust solutions based on Voronoi tessellations, grid maps [21], also integrating high-level features such as semantic knowledge [22] or communication channels [23] are available. Nevertheless, some works dealt with extracting or optimizing a graph directly taking into account the patrolling task [24–26]. Recent approaches are focusing on this problem from an on-line point of view where a graph representations is built starting from an unknown environment. An example is reported in [27] where a method based on structured triangulation is employed to build a representation for exploration and patrolling with a number of low-cost robots. In [14] the patrolling graph, embedded in a 3D space, is computed either from manually placed way-points or from

a history of robot trajectories. Finally notice that this type of discrete representation is clearly not the only one adopted in the literature. As will be mentioned later, some specific works rely on different continuous formulations.

Features. Each vertex can be associated to a feature vector where the interpretation of each element depends on the particular setting. Widely adopted features for vertices are the following.

- The *value*, denoted as $v_i \geq 0$, defines the importance of a vertex $v_i \in V$. Sometimes also called priority, it encodes the rationale that more important vertices correspond to more critical regions of the environment where a security breach would entail a large cost. The value entails the existence of a subset of targets $T = \{v_i \in V \mid v_i > 0\}$ that are interpreted as locations where a threat might occur. This feature can also be used as a rate at which the urge to patrol a vertex increases while it is left unguarded [28].
- The *patrolling cost* $c_i \geq 0$ quantifies the effort to assess the security status of vertex v_i . Examples can be found in situations when, once arrived at a vertex, one or more robots have to scan the location with their sensors or physically clear the area from the presence of possible intruders/hiders. In some settings this element is modeled as part of the solution, requiring to find the right amount of patrolling effort to allocate on each target in time [29].
- The *strength* $\delta_i > 0$ for a target $t_i \in T$ provides a measure of the effort needed to compromise the security status of vertex t_i . A large δ_i could represent strong resistance of target t_i to an attack, resulting in extended opportunities for detection by a patrolling robot. This feature is typically given as an amount of time needed to complete a breach on a target (indeed it is often called penetration time or attack time). In such cases, it can entail a minimum patrolling frequency constraint for each target (or equivalently a maximum lag between subsequent patrols) to guarantee full protection of the environment [15, 30].

Robots/agents. The team of patrolling robots/agents can be defined with the set $R = \{r_1, r_2, \dots, r_m\}$. In some cases this element is left implicit or just given with $m \in \mathbb{N}_{>0}$, in other cases is not given as a part of the problem and must instead be computed as a part of the solution [15, 31, 32].

Movement. A set $A = \{(v_i, v_j) \mid v_i \in V, v_j \in V \cup \{\perp\}\}$ lists the connections among ordered pairs of vertices of the graph; these are often in correspondence with the possible movements a robot can take from a currently occupied vertex v_i . The arc (v_i, v_j) (with $v_i \neq v_j$) represents the opportunity to travel from v_i to v_j without visiting any other vertex on the way to the destination; (v_i, v_i) , instead, can be used to model situations where a robot *leaves* v_i and then returns to it without patrolling other vertices; finally, (v_i, \perp) can represent the situation in which a robot stays at v_i without leaving

it. Each $a \in A$ is typically associated with a traveling cost $\omega_a \geq 0$ (usually a time or a distance) with the set of costs admitting a metric layout of the graph. Richer models can expand the movements set and the associated costs with additional features describing robots' locomotion models, for example accounting for orientation of movement and the cost of changing it [33].

Communication. Robots might need to communicate between each other in order to undertake some form of on-line coordination. The availability of communication channels can be included in the model with an additional set of arcs $C = \{(v_i, v_j) \mid v_i, v_j \in V\}$, where (v_i, v_j) encodes the possibility for two robots to exchange information when being at vertices v_i and v_j , respectively. The subgraph induced by C is sometimes called connectivity graph [32, 34–37].

Perception. The basic formulation assumes that the detection of a threat takes place as soon as any robot visits the vertex where the malicious activity is undergoing. General perception models can be described with a visibility function $s : V \rightarrow \mathcal{P}(V)$ where $s(v_i)$ is the subset of vertices that can be sensed from vertex v_i (usually including v_i itself) [29, 37]. In some scenarios robots' increased perceptions can be characterized by spatial uncertainty [38] or by detection errors [39, 40].

An abstract instance of a patrolling problem inspired by the above modeling ingredients is depicted and briefly described in Fig. 1.




Fig. 1 A graph-based patrolling setting composed of 5 targets. Assume, for simplicity, that symmetric movements are always possible, traveling costs and values are unitary, while $\delta_i = 2$ for each i . Two robots, r_i and r_j , follow two paths starting their mission at time $t = 0$ from target v_2 and v_5 , respectively. If patrolling costs are set to 0 and perception is limited to the currently occupied target, the idleness profile at $t = 4$ is $(3, 1, 2, 1, 2)$. In this scenario, an attack to target v_2 at $t = 1$ (when r_i is at v_1 and r_j is at v_3) would be successful. An attack on v_4 at $t = 2$ would fail. If patrolling costs are set to 2, the idleness profile at $t = 4$ becomes $(1, 3, 0, 5, 0)$, where a value of 0 indicates the presence of a robot on the corresponding vertex. If the same graph is adopted to model connectivity, robots can always exchange information except during their last visit, namely when occupying v_4 and v_2 , respectively

A patrolling setting is typically complemented with a threat-generating process which is characterized by precise assumptions and that, for the sake of simplicity, we shall refer to as *attacker*. All methods to compute patrolling strategies consider, to some extent, the presence of an attacker. A paramount distinction, however, is generally made on how the attacker is modeled. For a first class of works the attacker is *implicit* and indirectly modeled through a single/multi-objective function that the patrolling strategy has to optimize. Typical approaches adopt metrics that consider how the robots' visits to the various parts of the environment are distributed in space and time, to ensure that critical spots are patrolled enough often. The most popular concept to define such criteria is the *idleness*, namely the temporal delay between subsequent visits to a target [12, 13, 41]. Worst-case and average idleness are typical examples of adopted objective functions with many variants, substantially related, studied in literature. The cost of the patrolling routes (typically a traveling time or distance) can also be taken into account in order to achieve some level of efficiency. Costs are tightly connected to the patrolling performance, since more efficient strategies typically result in a better usage of resources and hence in an enhanced protection of the environment. These family of methods is often classified under the term of *non-adversarial* patrolling.

A second class of works considers *explicit* attacker models that generate threats according to given mechanisms. These attackers range from simple random processes or heuristics based on the observations of the past patrolling visits to more sophisticated rational fully or partially informed decision makers for which game-theoretical methods are most suitable [16–18, 19•, 20, 42, 43, 44•, 45–47]. These methods typically exploit non-deterministic solutions and try to maximize the expected probability of detecting attacks [17, 37, 48••]. This class of works is often referred to as *adversarial patrolling*.

To some extent, the research conducted in the last years showed how the distinction between adversarial and non-adversarial methods might be less sharp than expected. Key theoretical results in the non-adversarial domain provide insights on how to guarantee performance also in the presence of rational attackers. On the other side, adversarial methods often suffer from idealistic assumptions about the attacker which need to be refined to cope with less rational and more limited threat-generation capabilities.

Recent Trends

Today's research on robotic patrolling follows approaches that can be seen as in steady continuity with some of the challenges and methods that were introduced by the first seminal works. One of the papers introducing the patrolling

problem to the multi-agent community was proposed in [49]. The author provides a first formulation of the problem in multi-agent terms, also introducing *idleness* as an optimization criterion for good strategies. The paper discusses a first set of methods, focusing on the agent-based architecture that should implement the patrolling task, defining types of coordination and information sharing mechanisms. The work is complemented by [50] where a more theoretical perspective of the problem is taken by highlighting its close relation to the Traveling Salesman Problem [51]. A well-representative overview of these early and founding contributions can be found in [6], paving the way on how methods based on optimization and learning can be adopted for the problem. These works were among the first in highlighting how the robotic patrolling problem has deep connections with different combinatorial route-finding problems. This relation is today widely recognized and robotic patrolling is often cited among the representative applications domains of traveling salesman problems and their variants [52, 53].

Optimal Patrolling

A great deal of theoretical and algorithmic results for patrolling strategies that optimize some objective function can be found in a well-established research line in the automatic control community, which deals with a problem that shares a great number of basic features with robotic patrolling. The problem is customarily termed as *persistent monitoring* [54–56] since it takes the more general perspective of repeated information gathering on a set of targets without necessarily interpreting such information as finalized to the detection of a threat, which instead is typical in robotic patrolling. To some extent, persistent monitoring can be seen as a generalized formulation of non-adversarial patrolling. A distinctive feature of this line of works is that of casting the problem of computing patrolling strategies in an optimal control framework, hence providing resolution approaches that could go beyond discrete graph-based representations or networks. Optimality is often defined with respect to objective functions based on idleness or generalizations of it. An interesting formulation of this problem, for example, uses target-specific uncertainty costs growing at some rate (see values as discussed in Section 2) when no agent is inspecting the target and decreasing at some other rate when one or more agents are instead monitoring it. The target dwelling time, namely the time spent monitoring a particular target, is modeled as a part of the solution and not assumed as fixed in advance (see patrolling costs defined in Section 2). As recently pointed out in [29], this formulation has intriguing connections with queuing theory models, since the increase/decrease of the uncertainty cost can be interpreted as arrival/servicing operations on the target while dwelling time can be seen as the amount of service allocated to a target at a particular time.

A recent example of a problem formulation of the above kind has been proposed in [39] for settings defined on one-dimensional spaces over a finite horizon with agents characterized by a probabilistic sensing model. The problem is solved by computing agent trajectories considering second-order dynamics (accelerations) and by tackling the optimization problem with a gradient-based method. The extension of this type of approaches to two-dimensional spaces could provide key advancements for robotic patrolling. Another example is investigated in [28] where authors assume that costs accumulate at vertices with a growth rate described by a probabilistic model. The proposed approach estimates the cost levels and tackles the problem as a variant of the team-orienteeing problem.

The optimality of patrolling can clearly be described in terms of minimizing the worst-case idleness on a graph. The work proposed in [12] defines approximation methods for this problem as well as an exact algorithm for a special 1D case. In [13] a richer setting is considered where the patrolling robot has limited autonomy and needs to return at a service station to be recharged after a maximum number of visits it can make on a graph-represented environment. Recharging has a fixed temporal cost called service time. Authors provide an algorithmic characterization of the optimal solution showing that, for some values of the service time, the general problem might either reduce to a finite subset of its instances or substantially become a TSP. For the remaining cases a method with an additive approximation guarantee (equal to the service time) is derived.

The hardness and algorithmic results of the above works from one side provide key further understanding for the problem of computing optimal cyclic patrolling routes in the presence of realistic constraints [57], but from the other promote methods capable of gaining computational efficiency at the expense of solution quality (examples include approximation or heuristic methods). A very well-known and widely applied result in robotics concerns sub-modular optimization: when the function describing a total information reward conquered by a sequence of observations delivers diminishing returns, greedy methods for building such a sequence tend to perform well also providing approximation guarantees [58]. This idea has been recently combined with receding horizon planning in [59] to synthesize patrolling strategies with bounded optimality gap. Another approach is to exploit decentralization. A recent example is proposed in [60] where a gradient-based method for a given objective function to be minimized is followed.

Optimal patrolling strategies could also be defined in terms of satisfaction of idleness constraints. The works proposed in [15, 61] provide key complexity results for the problem of patrolling over a graph under deadline constraints, substantially defined by a set of maximum idleness values, one for each target. This problem formulation has

implications both in the non-adversarial and adversarial domains of patrolling. In the first case deadlines can be interpreted as efficiency requirements (alternative formulations of this problem have been investigated in terms of frequency constraints, also accounting for on-line events affecting the patrolling activities [62]). In the second case deadlines could just coincide with targets' strength (recall Section 2), and solving the aforementioned problem would entail the detection of any possible attack. In [30] authors propose an approximation algorithm and a heuristic method for such a problem. Experimental comparisons show how the heuristics (based on the orienteering problem) usually perform better.

On-line Cooperation for Dynamic Settings

Robustness and adaptability to external factors that can hinder the execution of patrolling missions in the real world are of key importance and pose the need for the patrolling strategy to timely react. An example of this need is provided in the scenario outlined in [11] where a dynamic topology is adopted by assuming that the availability of arcs (see movements as discussed in Section 2) over time is only known in probability.

Distributed algorithms exhibiting scalability and robustness to unforeseen events have been proposed [63]. A recent line of research dealt with this problem by considering an on-line setting where interference among robots, triggered obstacle avoidance behaviors, and communication delays might be among those events [41]. The patrolling problem is addressed from the perspective of on-line task-assignment where rewards are based on idleness and that is solved in a decentralized way. A greedy task-assignment method and another one based on sequential single-item auctions are proposed and evaluated confirming how on-line coordination can be used to achieve good performance in real dynamic scenarios. In [34] time-variance is modeled for a notion of status applied to each vertex. A state-information variable is associated to each vertex to describe the current (discrete) level of potential gain the patroller can get by visiting it. This level changes according to an unknown transition model. The problem is solved in a decentralized fashion by allowing agents to make local observations and cooperate within a Bayesian reinforcement learning setting where on-line learning and planning are executed. The work in [29] proposes an on-line and distributed method that tries to improve on the classical gradient-based ones, typically suffering from local optima problems. The problem of where to patrol next is solved, in each agent, with an event-driven receding horizon control approach. Agents plans are computed upon observation of local events (like, for example, arrivals or departures in neighboring vertices) up to some automatically optimized horizon length. In [64] authors propose a method to compute

patrolling strategies that can integrate real-time data from the environment, including criminal events or emergencies. This additional information can change, at run-time, the priorities of targets and, as a consequence, the strategy must be adapted.

As some of these works show, cooperation by means of on-line information exchange might play a crucial role. Literature provides established methods to achieve cooperation by means of distributed intelligence, one significant example for patrolling being the use of Bayesian learning [65]. Recently, the work proposed in [14] addresses patrolling coordination in a unified framework that encompasses both high-level decisions on a patrolling graph and low-level resolution of navigation conflicts between robots in a 3D environment. The proposed distributed method seeks idleness optimization at the high level while performing path planning based on 3D SLAM and traversability analysis at the low level. In [35] robots communicate between neighboring regions to detect faults. Such interactions are exploited in a scenario considering idleness and isolation time, namely the expected time a robot stays without communicating with any other one.

Recently, emerging methods in the field of graph neural networks are beginning to be investigated. The work in [66] proposes to use multi-agent reinforcement learning to perform coordination between agents monitoring a graph. The approach is based on a graph attention network [67] to allow agent sharing of locally perceived information. Authors show how this method could indeed introduce some advantages with respect to classical patrolling baselines (greedy or TSP-based) providing evidence on how coordination through communication is valuable and suggesting how the method could reach, in the form of an emergent behavior, the adoption of cyclic periodic paths solving the underlying patrolling problem.

Indirect communication through the environment has also been investigated as a method to obtain cooperation among patrolling robots, a paradigm typically studied in swarm robotics [68]. An important characteristic of these solutions is to obtain emerging team coordination by means of simple and efficient individual behaviors. Some recent works are devoting efforts to apply this approach to patrolling. In [69] authors focus on simplifying agents by reducing their range of perception according to which they can collect the information left by others in the environment. Specifically, the work provides a method to convert strategies for agents that can perceive the current and neighboring vertices to ones that require to perceive only the current vertex. A key result is that reduced-range strategies exhibit competitive performance, suggesting how even very limited sensing capabilities might be used to obtain effective coordination for patrolling. In [70] authors propose a pheromone-based coordination to communicate uncertainty costs about the

targets to deal with the patrolling setting introduced in [71] where a patrolling robot with given sensing capabilities has to perform persistent road-network patrol in order to detect probabilistic events.

Cooperation among robots, which is enabled by the exchange of information between them, can be also enforced by the presence of communications or connectivity constraints. Communications during the patrolling mission can be required to report gathered information to a base station. In [36] authors study patrol routes that try to minimize the absolute time required to gather and report at the base station as well as the relative time between information discovery and its reporting. The proposed approach, based on the computation of minimum-latency paths, shows how cooperation between the robots in undertaking a store-and-forward behavior can increase the performances. The same authors consider in [72] the minimization of the communication lag together with visits-induced idleness.

The requirement of staying connected can be instantiated in many different forms and is a central problem in many robotics scenarios including robotic patrolling [73]. The presence of these requirements for a team of robots is addressed in [32]. Robots must patrol an environment while always forming a connected component in order to maintain multi-hop connectivity with a base station. To assess when, at any given time, two robots are connected to each other a set of connectivity edges is defined over the graph that models the environment. Authors provide key NP-Hardness results for the problem of minimizing the worst idleness induced on the graph, finding the minimum number of robots to ensure each target can be reached while maintaining connectivity, and optimally placing relay robots to support communications. Tree-traversal, partitioning, and heuristic methods for convex grids are proposed to obtain patrolling strategies that could guarantee good performance while complying with the connectivity requirements.

Dealing with Attackers

When explicitly including an attacker in the patrolling problem, it is commonly assumed that this one can gain, by performing environmental observations, some degree of knowledge about the patrolling strategy. Such a knowledge can then be exploited to maximize the chances of a successful attack, for example by predicting with some level of confidence the next moves of the patroller. Knowledge can be about the strategy, i.e., the law the patrollers are using to decide where to patrol next and/or its realization, namely the history of patrols performed up to the current time. One way to counteract this issue is to adopt stochastic patrolling strategies, where the movements are drawn from some probability distribution. The approach presented in [37] computes non-deterministic patrolling strategies as

Markov chains minimizing expected back-and-forth traveling times among pairs of vertices that are induced by the transition probabilities. The work combines in such a setting interesting realistic features like limited sensing for the robot, need to communicate with a base station in a centralized fashion, and evaluate the probability of capturing an attacker by estimating the amount of time vertices are left unguarded. The work proposed in [74] deals with this problem by considering probabilistic inspection constraints for the patrolling routes, namely requirements on the expected number of agents patrolling each area at a random time. The problem is formulated with signomial programming and its resolution is tackled with a distributed method. A significant line of research deals with this problem in patrolling perimeters. In [33] environments are modeled with open polylines divided in segments that can be traversed by a robot traveling between endpoints. In this formulation, targets are represented by edges connecting two endpoints, which an adversary with full knowledge of the patrollers' strategy might try to penetrate. The patrolling strategy is encoded, at each target, with a probability of changing travel direction for a passing-by robot. The paper improves on previous solutions adopting equal probabilities for each target and shows how cooperation schemes in such a setting might further improve the performance.

The most studied and adopted model for stochastic patrolling strategies is perhaps given by Markov chains, especially first-order ones that condition the next move upon the currently occupied vertex. A significant analysis of these models is presented in [75], where authors formulate optimization problems of the patrolling problem for different scenarios. The notion of hitting times is leveraged as a proxy for the speed of a patrolling strategy. Also the concepts of entropy rate and entropy of the return times of a patrolling strategy are analyzed with respect to its implication in adversarial scenarios. Strategies maximizing these metrics tend to leak less information to an adversarial observer.

Modeling the observation process that an attacker might carry out is another recent trend studied in robotic patrolling. In [16, 17] the idea of an attacker with limited observation capabilities was proposed. Specifically, the attacker is assumed of being capable of (imperfectly) observing the presence of the patroller only locally to a single target. No other knowledge, for example, about the graph topology or the patrolling strategy, is assumed to be accessible. The work proposes a set of heuristic methods to exploit this limitation to the advantage of a patroller. Solutions are focused on performing patrolling while revealing the least information to the adversarial observer by means of strategic injections of delays along the patrolling routes and the adoption of time-variant Markov chains. In [18] authors consider, in a graph-patrolling setting, an attacker with a limited time to observe the patrolling strategy's realization on the graph.

The decision about whether to attack and, in the positive case, when and on which target accounts for the confidence of the knowledge about the patrolling strategy extracted from the observations.

Game theory is a natural mathematical framework for modeling and, more importantly, characterizing the solution (in terms of equilibrium points) of the adversarial interaction that such models entail. The capability of attackers to observe has been considered also by recent contributions in this scope. A significant example is [42] where the problem of patrolling a perimeter is analyzed. In the proposed formulation, robots are dispatched from a base at a given rate, each covering a cycle along chosen direction and speed (time and space are continuous). Attacks can take place at any point on the perimeter and can be detected by any passing-by patroller with some given probability. The attacker is allowed to observe the activity of the patrollers at the chosen target to better inform her strategy. The paper provides key theoretical properties on the zero-sum (also Stackelberg, see below) game resulting from this setting, also showing how an optimal detection probability can be achieved by the patrollers whether they are observable or not. In [19•] authors consider a game played on a star-graph where targets are placed on endpoints. The patroller applies a non-deterministic strategy described by a Markov chain while the attacker secretly chooses a target and conditions its decision to perform an attack to the history of presence/absence of the patroller on such a target. The paper deals with the problem of determining the optimal strategy (maximizing the capture probability) and hence providing the value of the zero-sum game. A key insight is provided about deterrence, namely the capability of discouraging attacks by exposing the realization of the patrolling strategy to the observing attacker (a feature that seems to be deeply connected to the non-deterministic nature of the patrolling strategy).

To some extent, deterrence can be seen as a specific case of deception, a method that for perimeter settings has been studied in [43]. Such a work studies methods to mislead an attacker that is not assumed to have full knowledge of the patrolling strategy and its execution. Instead it can be manipulated, by inducing particular observations of the robots, into believing a stronger patrolling profile with respect to the actual one. The work shows how to perform two types of deception exploiting either limited observation capabilities of the attacker or the use of dummy robots (cheaper units without detection capabilities). The results highlight a trade-off between guaranteeing that the attacker does not discover the deception attempt and the resulting capture probability.

When adversarial observing capabilities are taken from a worst-case stance, a Stackelberg game is typically obtained. This setting sees a defender (also known as leader, the agent controlling one or more patrolling robots) acting under the assumption that the attacker (follower) has perfect

knowledge of the patrolling strategy to which she will best respond. Robotic patrolling settings following this approach can be seen as a subset of a more general class of game-theoretical models called Stackelberg (or Leader-Follower) Security Games (SSGs), which has a longstanding record both in terms of research and real-world applications [76].

The SSGs formulations proposed for patrolling [38, 77], originally defining heuristic algorithmic methods, have been recently subject of deeper mathematical analysis. Notably, the work proposed in [48••] provides key theoretical results in deriving a universal upper bound on the performance obtainable in such games, defining optimal strategies for linear and star graphs, and approximated ones for complete graphs. Recent contributions in this domain considered spatio-temporal constraints for the realization of the patrolling strategy [78], explored the use of Monte Carlo Tree Search as resolution method [10, 79], integrated data-driven learning of the adversarial preferences and behavior [80], and exploited reinforcement learning to deal with incomplete state information [81].

Classical game-theoretical formulations, not considering a Stackelberg paradigm, have also been studied for patrolling with potential applications to robots. A seminal work that considered the problem in such terms is [82]. In recent contributions, proposed in [20], authors consider a two-player game formulation played on a graph between a patroller periodically covering a path and an attacker choosing a specific vertex and time for the attack. The work proposes a characterization for the game's value, defined as the probability of capturing the attacker, providing optimal strategies in some specific cases (linear graphs, and arbitrary graphs with even periods for the patrolling strategy). Another example can be found in [83], where the patrolling game is settled on a continuous space.

The attacking dynamics encoded in these models have also been subject of enhanced modeling, in the attempt of capturing more realistic or more challenging attacker behaviors. In [44•], for example, authors enrich the patrolling problem with the need to respond to attacks instead of merely detect them. This means that a detection could, in principle, trigger additional actions for the robots to undertake. Such actions might alter the patrolling schedule and generate vulnerabilities at run-time. The paper provides an approach to compute optimal patrolling strategies in a setting where an attacker might strategically perform sequential attacks, with the attempt of triggering and exploiting vulnerabilities via patrolling responses. Another variant of the attacker model has also been proposed in [45] where authors consider attackers that can choose the duration of the attack to carry out. In [46] attackers can undertake deception to hide their preferences.

As many of the above works suggest, the use of Markov chains still today represents the mainstream approach for

non-deterministic patrolling strategies, mainly due to their simplicity and ease of implementation. At the same time, their limits are very well-known, especially for the widely adopted setting where the next patroller move only depends on the currently occupied position. A recent alternative approach has been discovered in [84, 85] with what the authors call regular strategies. These are methods where the patroller's decision on where to go next still depends on finite information computed from the history of previous patrols, but can provide comparative performance with respect to Markov strategies that condition the next move on the whole history of visits (hence being more powerful but also intractable). The most recent results of this line of works provide methods to efficiently synthesize these strategies in realistic graph-based settings.

Another of the drawbacks of patrolling strategies expressed with Markov chains is that computing the detection probability for a potential attack is a difficult task that might require costly numerical simulations or the resolution of non-linear systems of equations [75, 77]. One recent advancement in this scope has been proposed in [47]. Authors deal with patrolling on polyline graphs against a full-knowledge attacker and proposed a combinatorial method based on lattice paths. The method allows for efficiently computing the number of possible patrolling paths to derive an expression for the probability of capture.

Evaluating Performances and Costs

The definition of widely adoptable methods for the empirical evaluation of the performance of patrolling strategies is a central problem, especially due to the fact that deploying and testing the above methods in the physical world, using real robots, has typically large costs. Understandably, this is still a largely unaddressed task with only a minority of works focusing on empirical analysis of a real-world implementation [14, 28, 37, 41, 63, 69]. One emerging way to lower the barrier to this type of analysis is the use of realistic simulators, as they represent a well-established approach in many robotics applications to decrease the costs of experimentation [31]. A ROS-based simulation framework for patrolling has recently been proposed in [86], where new algorithms for patrolling strategies can be integrated and compared abstracting away from their low-level implementation on robots. The tool allows efficient testing for instructing real-world subsequent experimental campaigns.

The number of deployed robots is a particularly cost-sensitive dimension when dealing with real scenarios. In [87] authors deal with the problem of estimating the performance, in terms of an idleness-based objective, of a given number of robots employed in patrolling missions with the aim of dimensioning their number with respect to a required performance level.

Future Directions

The general picture that recent research on robotic patrolling depicts is a very rich and heterogeneous one. Researchers from different communities dealt with this problem by means of different approaches and methods. One shared trend, however, seems to emerge as a growing interest towards formulations with additional descriptive power with respect to prior work, in the attempt of better adhering to the reality at stake. Examples are the inclusion of constraints (involving sensing, energy, or communication), the need to handle on-line events, the adoption of limited adversary models, and the inclusion of additional decision variables. Around this general trend some interesting future research directions can be envisaged.

Instance extraction. A very practical, but at the same time crucial, issue is how to generate model instances for specific use cases at hand. As expressiveness improves with richer formulations, the number of parameters required to fit a suitable representation of the environment increases. Prior works suggest that extracting the right representation might have an impact on the patrolling performance whether the problem domain is considered as adversarial [26] or not [24]. Concrete examples of this issue have been recently addressed in [27] and [14] concerning the representation of the environments and in [13] to determine robots' maximum operational range induced by energy limits. The specificity and sparseness of such solutions pose the need of a more general methodology where metrics like sensitivity and scalability could be assessed and, ideally, evaluated with benchmarks.

Tailored models. Specific use cases might admit tailored formulations and resolution methods where efficiency and performance might be improved at the cost a reduced generality. An example of this approach can be found in [88]. In such a work, authors deal with a patrolling formulation inspired by a maritime scenario. They adopt a multi-objective formulation combining the total sum of distances traveled by robots, the maximum distance traveled by a robot, and the priority of targets. A bio-inspired method adopting immune and endocrine system dynamics is exploited for the problem resolution. The resulting strategies exhibit interesting performance when compared against heuristic methods for the multi-objective problem. Similarly specific approaches might emerge for other patrolling use cases.

Long-term reliability. The study of issues related to long-term autonomy is a topic of growing relevance for many robotic application domains, including patrolling [89]. Indeed, robotic patrolling is often meant to be a persistent and unsupervised activity where robots have to be operational for long time spans. This scenario opens for additional challenges since reliability can face increased risks:

patrolling strategies might need to cope with failures or vulnerabilities difficult to forecast or reproduce in simulated/real campaigns that are necessarily carried out in short/mid-term setups. A recent example work dealing with a related issue has been proposed in [90]. Authors consider the problem of patrolling an environment with an autonomous robot and focus on providing robust and reliable localization.

Integration An interesting direction of development can also be represented by integrated patrolling settings. The idea can be interpreted with two different perspectives from the point of view of a multi-robot system. An outward integration implies a richer real-time exchange of information between the robots and the environment. The detection of threats, hence, could not only be based on robots' local perceptions, but also to other data that the environment allows to access on a global frame. Examples can be found in patrolling for industrial environments where live indicators on, for example, the current workload, the efficiency of the production processes, or the presence of faults might be accessible. Recent efforts are applying advanced deep learning methods, such as variational autoencoders, to perform anomaly detection in these settings [91]. The integration of such techniques in robotic patrolling can endow patrolling strategies with the capabilities of counteracting more sophisticated threats. On the other side, inward integration is meant to increase the level of introspection of a multi-robot system by means of patrolling capabilities. Similarly to what recent works are starting to consider for multi-agent path planning problems [92], patrolling methods can play a role in the presence of self-monitoring requirements where surveillance applies to the robots themselves during the execution of a different task.

Defenders and attackers On the adversarial side, a rigorous mathematical study of more sophisticated game models and algorithmic solutions is surely a key direction of research whose pursuing is likely to deepen the understanding of how the optimality of patrolling strategies manifests under strategic terms. Observation capabilities have been studied in different flavors from the attacker point of view. With the exception of some works in the field of security games [80], this feature has been less studied from the point of view of the patrolling robots. One emerging approach is to incorporate the attacking dynamics into a reward function and exploit a reinforcement learning framework to synthesize a patrolling strategy by means of multiple interactions with the environment. A similar idea was originally proposed in [93] and has been recently reconsidered by [94, 95] in the context of deep reinforcement learning. Further development of these methods can

provide strategies that can learn to deal with arbitrary attackers, eliminating the need of formalizing attacking behaviors in detail. Moreover, modeling deceptive behaviors inside this approach represents an intriguing direction of investigation [96].

Conclusions

Robotic patrolling is a problem investigated in different research communities with the common goal of realizing autonomous on-the-field surveillance systems based on mobile robots. Researchers have spent a considerable effort during the past decades in defining effective methods for computing patrolling strategies, a key component by which robots can allocate their surveillance activities in space and time to achieve protection of a physical environment. This paper provides a review of some of the significant advancements that characterized recent research, with a particular attention to works exhibiting theoretical or practical relevance within the field of mobile robotics and multi-agent systems. The analysis of such works shows how this field is still very active and heterogeneous, with a common trend of targeting more realistic and complex formulations in full accordance with the objectives and scenarios envisaged by the early seminal works proposed in the literature.

Funding Open access funding provided by Università degli Studi di Milano within the CRUI-CARE Agreement.

Declarations

Conflict of Interest The author declares no competing interests.

Human and Animal Rights and Informed Consent This article does not contain any studies with human or animal subjects performed by any of the authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Papers of particular interest, published recently, have been highlighted as:

- Of importance
- Of major importance

1. Rubio F, Valero F, Llopis-Albert C. A review of mobile robots: Concepts, methods, theoretical framework, and applications. *International Journal of Advanced Robotic Systems*. 2019;16(2):1729881419839596.
2. Alatise MB, Hancke GP. A review on challenges of autonomous mobile robot and sensor fusion methods. *IEEE Access*. 2020;8:39830–46.
3. Rizk Y, Awad M, Tunstel EW. Cooperative heterogeneous multi-robot systems: a survey. *ACM Computing Surveys (CSUR)*. 2019;52(2):1–31.
4. Quattrini Li A. Exploration and mapping with groups of robots: Recent trends. *Current Robotics Reports*. 2020;1(4):227–37.
5. Queralta JP, Taipalmaa J, Pullinen BC, Sarker VK, Gia TN, Tenhunen H, et al. Collaborative multi-robot search and rescue: Planning, coordination, perception, and active vision. *Ieee Access*. 2020;8:191617–43.
6. Almeida A, Ramalho G, Santana H, Tedesco P, Menezes T, Corruble V, et al. Recent advances on multi-agent patrolling. In: *Brazilian Symposium on Artificial Intelligence*. Springer; 2004. p. 474–83.
7. Portugal D, Rocha R. A survey on multi-robot patrolling algorithms. In: *Doctoral conference on computing, electrical and industrial systems*. Springer; 2011. p. 139–46.
8. Huang L, Zhou M, Hao K, Hou E. A survey of multi-robot regular and adversarial patrolling. *IEEE/CAA Journal of Automatica Sinica*. 2019;6(4):894–903.
9. Samanta S, Sen G, Ghosh SK. A literature review on police patrolling problems. *Annals of Operations Research*. 2021;1–44.
10. Kartal B, Godoy J, Karamouzas I, Guy SJ. Stochastic tree search with useful cycles for patrolling problems. In: *IEEE international conference on robotics and automation (ICRA)*. IEEE. 2015;2015:1289–94.
11. Othmani-Guibourg M, El Fallah-Seghrouchni A, Farges JL, Potop-Butucaru M. Multi-agent patrolling in dynamic environments. In: *IEEE international conference on agents (ICA)*. IEEE. 2017;2017:72–7.
12. Afshani P, Berg Md, Buchin K, Gao J, Löffler M, Nayyeri A, et al. Approximation algorithms for multi-robot patrol-scheduling with min-max latency. In: *International Workshop on the Algorithmic Foundations of Robotics*. Springer; 2020. p. 107–23.
13. Hari SKK, Rathinam S, Darbha S, Kalyanam K, Manyam SG, Casbeer D. Optimal UAV route planning for persistent monitoring missions. *IEEE Transactions on Robotics*. 2020;37(2):550–66.
14. Freda L, Gianni M, Pirri F, Gawel A, Dubé R, Siegwart R, et al. 3D multi-robot patrolling with a two-level coordination strategy. *Autonomous Robots*. 2019;43(7):1747–79.
15. Drucker N, Ho HM, Ouaknine J, Penn M, Strichman O. Cyclic-routing of unmanned aerial vehicles. *Journal of Computer and System Sciences*. 2019;103:18–45.
16. Basilico N, Carpin S. Balancing unpredictability and coverage in adversarial patrolling settings. In: *International Workshop on the Algorithmic Foundations of Robotics*. Springer; 2018. p. 762–77.
17. Alvarenga CD, Basilico N, Carpin S. Time-varying graph patrolling against attackers with locally limited and imperfect observation models. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE; 2019. p. 4869–76.
18. Asghar AB, Smith SL. A patrolling game for adversaries with limited observation time. In: *2018 IEEE Conference on Decision and Control (CDC)*. IEEE; 2018. p. 3305–10.
19. ●Alpern S, Chleboun P, Katsikas S, Lin KY. Adversarial Patrolling in a Uniform. *Operations Research*. 2022;70(1):129–40. **This paper derives key theoretical results for a patrolling game where the concept of deterrence is analyzed.**
20. Alpern S, Lidbetter T, Papadaki K. Optimizing periodic patrols against short attacks on the line and other networks. *European Journal of Operational Research*. 2019;273(3):1065–73.
21. Lau B, Sprunk C, Burgard W. Improved updating of Euclidean distance maps and Voronoi diagrams. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE; 2010. p. 281–6.
22. Kostavelis I, Gasteratos A. Semantic mapping for mobile robotics tasks: A survey. *Robotics and Autonomous Systems*. 2015;66:86–103.
23. Quattrini Li A, Penumarthy PK, Banfi J, Basilico N, O’Kane JM, Rekleitis I, et al. Multi-robot online sensing strategies for the construction of communication maps. *Autonomous Robots*. 2020;44(3):299–319.
24. Kolling A, Carpin S. Extracting surveillance graphs from robot maps. In: *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE; 2008. p. 2323–8.
25. Portugal D, Rocha RP. Retrieving topological information for mobile robots provided with grid maps. In: *International Conference on Agents and Artificial Intelligence*. Springer; 2012. p. 204–17.
26. Basilico N, Gatti N. Automated abstractions for patrolling security games. In: *Twenty-Fifth AAAI Conference on Artificial Intelligence*; 2011.
27. Lee SK, Fekete SP, McLurkin J. Structured triangulation in multi-robot systems: Coverage, patrolling, Voronoi partitions, and geodesic centers. *The International Journal of Robotics Research*. 2016;35(10):1234–60.
28. Liu B, Xiao X, Stone P. Team orienteering coverage planning with uncertain reward. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE; 2021. p. 9728–33.
29. Welikala S, Cassandras CG. Event-driven receding horizon control for distributed persistent monitoring in network systems. *Automatica*. 2021;127: 109519.
30. Asghar AB, Smith SL, Sundaram S. Multi-robot routing for persistent monitoring with latency constraints. In: *American Control Conference (ACC)*. IEEE. 2019;2019:2620–5.
31. Collins J, Chand S, Vanderkop A, Howard D. A review of physics simulators for robotic applications. *IEEE Access*. 2021.
32. Scherer J, Rinner B. Multi-robot persistent surveillance with connectivity constraints. *IEEE Access*. 2020;8:15093–109.
33. Oshrat Y, Agmon N, Kraus S. Adversarial Fence Patrolling: Non-Uniform Policies for Asymmetric Environments. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34; 2020. p. 10377–84.
34. Zhou X, Wang W, Wang T, Lei Y, Zhong F. Bayesian reinforcement learning for multi-robot decentralized patrolling in uncertain environments. *IEEE Transactions on Vehicular Technology*. 2019;68(12):11691–703.
35. Caraballo LE, Díaz-Báñez JM, Fabila-Monroy R, Hidalgo-Toscano C. Stochastic strategies for patrolling a terrain with a synchronized multi-robot system. *European Journal of Operational Research*. 2021.

36. Scherer J, Rinner B. Multi-UAV surveillance with minimum information idleness and latency constraints. *IEEE Robotics and Automation Letters*. 2020;5(3):4812–9.
37. Alam T, Rahman M, Carrillo P, Bobadilla L, Rapp B, et al. Stochastic multi-robot patrolling with limited visibility. *Journal of Intelligent & Robotic Systems*. 2020;97(2):411–29.
38. Basilico N, De Nittis G, Gatti N. Adversarial patrolling with spatially uncertain alarm signals. *Artificial Intelligence*. 2017;246:220–57.
39. Wang YW, Wei YW, Liu XK, Zhou N, Cassandras CG. Optimal persistent monitoring using second-order agents with physical constraints. *IEEE Transactions on Automatic Control*. 2018;64(8):3239–52.
40. Bondi E, Oh H, Xu H, Fang F, Dilkina B, Tambe M. To signal or not to signal: Exploiting uncertain real-time information in signaling games for security and sustainability. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34; 2020. p. 1369–77.
41. Farinelli A, Iocchi L, Nardi D. Distributed on-line dynamic task assignment for multi-robot patrolling. *Autonomous Robots*. 2017;41(6):1321–45.
42. Lin KY. Optimal patrol of a perimeter. *Operations Research*. 2021.
43. Talmor N, Agmon N. On the Power and Limitations of Deception in Multi-Robot Adversarial Patrolling. In: *IJCAI*; 2017. p. 430–6.
44. •Lin ES, Agmon N, Kraus S. Multi-robot adversarial patrolling: Handling sequential attacks. *Artificial Intelligence*. 2019;274:1–25. **This paper provides a key advancement in adversarial patrolling by considering an attacker capable of multiple, sequential attacking actions.**
45. Yang HT, Tsai SY, Liu KS, Lin S, Gao J. Patrol scheduling against adversaries with varying attack durations. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*; 2019. p. 1179–88.
46. Nguyen TH, Butler A, Xu H. Tackling Imitative Attacker Deception in Repeated Bayesian Stackelberg Security Games. In: *ECAI*; 2020. p. 187–94.
47. Buermann J, Zhang J. Multi-Robot adversarial patrolling strategies via lattice paths. In: *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*; 2021. p. 4213–9.
48. ••Duan X, Paccagnan D, Bullo F. Stochastic Strategies for Robotic Surveillance as Stackelberg Games. *IEEE Transactions on Control of Network Systems*. 2021;8(2):769–80. **This paper provides a universal upper bound on the performance obtainable in Stackelberg patrolling games. Optimal strategies for linear and star graphs are derived.**
49. Machado A, Ramalho G, Zucker JD, Drogoul A. Multi-agent patrolling: An empirical analysis of alternative architectures. In: *International workshop on multi-agent systems and agent-based simulation*. Springer; 2002. p. 155–70.
50. Chevaleyre Y. Theoretical analysis of the multi-agent patrolling problem. In: *Proceedings. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2004. (IAT 2004)*. IEEE; 2004. p. 302–8.
51. Feillet D, Dejax P, Gendreau M. Traveling salesman problems with profits. *Transportation science*. 2005;39(2):188–205.
52. Agatz N, Bouman P, Schmidt M. Optimization approaches for the traveling salesman problem with drone. *Transportation Science*. 2018;52(4):965–81.
53. Cheikhrouhou O, Khoufi I. A comprehensive survey on the Multiple Traveling Salesman Problem: Applications, approaches and taxonomy. *Computer Science Review*. 2021;40: 100369.
54. Song C, Liu L, Feng G, Xu S. Optimal control for multi-agent persistent monitoring. *Automatica*. 2014;50(6):1663–8.
55. Yu J, Schwager M, Rus D. Correlated orienteering problem and its application to persistent monitoring tasks. *IEEE Transactions on Robotics*. 2016;32(5):1106–18.
56. Hari SKK, Rathinam S, Darbha S, Kalyanam K, Manyam SG, Casbeer D, The generalized persistent monitoring problem. In: *American Control Conference (ACC)*. IEEE. 2019;2019:2783–8.
57. Lauri F, Créput JC, Koukam A. The multi-agent patrolling problem theoretical results about cyclic strategies. In: *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer; 2014. p. 171–82.
58. Krause A, Guestrin C. Near-optimal observation selection using submodular functions. In: *AAAI*. vol. 7; 2007. p. 1650–4.
59. Rezazadeh N, Kia SS. A sub-modular receding horizon solution for mobile multi-agent persistent monitoring. *Automatica*. 2021;127: 109460.
60. Zhou N, Cassandras CG, Yu X, Optimal Andersson SB., Policies Threshold-Based Distributed Control, for Persistent Monitoring on Graphs. In: *American Control Conference (ACC)*. IEEE. 2019;2019:2030–5.
61. Ho HM, Ouaknine J. The cyclic-routing UAV problem is PSPACE-complete. In: *International Conference on Foundations of Software Science and Computation Structures*. Springer; 2015. p. 328–42.
62. Elmaliach Y, Agmon N, Kaminka GA. Multi-robot area patrol under frequency constraints. *Annals of Mathematics and Artificial Intelligence*. 2009;57(3):293–320.
63. Yan C, Zhang T. Multi-robot patrol: A distributed algorithm based on expected idleness. *International Journal of Advanced Robotic Systems*. 2016;13(6).
64. Rumi SK, Shao W, Salim FD. Realtime predictive patrolling and routing with mobility and emergency calls data. In: *Proceedings of the international AAAI conference on web and social media*. vol. 14; 2020. p. 964–8.
65. Portugal D, Rocha RP. Cooperative multi-robot patrol with Bayesian learning. *Autonomous Robots*. 2016;40(5):929–53.
66. Chen J, Baskaran A, Zhang Z, Tokekar P. Multi-Agent Reinforcement Learning for Visibility-based Persistent Monitoring. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE;. p. 2563–70.
67. Velickovic P, Cucurull G, Casanova A, Romero A, Lio P, Bengio Y. Graph attention networks stat. 2017;1050:20.
68. Chu HN, Glad A, Simonin O, Sempé F, Drogoul A, Charpillet F. Swarm approaches for the patrolling problem, information propagation vs. pheromone evaporation. In: *19th IEEE international conference on tools with artificial intelligence (ICTAI 2007)*. vol. 1. IEEE; 2007. p. 442–9.
69. Azevedo Sampaio P, da Silva Sousa R, Nazário Rocha A. Reducing the range of perception in multi-agent patrolling strategies. *Journal of Intelligent & Robotic Systems*. 2018;91(2):219–31.
70. Wang T, Dong G, Huang P. Pheromone-Diffusion-based Conscientious Reactive Path Planning for Road Network Persistent Surveillance. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE; 2021. p. 7922–8.
71. Wang T, Huang P, Dong G. Modeling and path planning for persistent surveillance by unmanned ground vehicle. *IEEE Transactions on Automation Science and Engineering*. 2020;18(4):1615–25.
72. Scherer J, Rinner B. Multi-robot patrolling with sensing idleness and data delay objectives. *Journal of Intelligent & Robotic Systems*. 2020;99(3):949–67.
73. Khateri K, Pourgholi M, Montazeri M, Sabattini L. A comparison between decentralized local and global methods for connectivity maintenance of multi-robot networks. *IEEE Robotics and Automation Letters*. 2019;4(2):633–40.

74. Hosseinalipour S, Rahmati A, Dai H, et al. Energy-aware stochastic UAV-assisted surveillance. *IEEE Transactions on Wireless Communications*. 2020;20(5):2820–37.
75. Duan X, Bullo F. Markov Chain-Based Stochastic Strategies for Robotic Surveillance. *Annual Review of Control, Robotics, and Autonomous Systems*. 2021;4:243–64.
76. Sinha A, Fang F, An B, Kiekintveld C, Tambe M. Stackelberg security games: Looking beyond a decade of success. *IJCAI*; 2018
77. Basilico N, Gatti N, Amigoni F. Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial intelligence*. 2012;184:78–123.
78. Clempner JB. A continuous-time Markov Stackelberg security game approach for reasoning about real patrol strategies. *International Journal of Control*. 2018;91(11):2494–510.
79. Karwowski J, Mańdziuk J. A Monte Carlo Tree Search approach to finding efficient patrolling schemes on graphs. *European Journal of Operational Research*. 2019;277(1):255–68.
80. Perrault A, Wilder B, Ewing E, Mate A, Dilkina B, Tambe M. End-to-end game-focused learning of adversary behavior in security games. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34; 2020. p. 1378–86.
81. Alcantara-Jiménez G, Clempner JB. Repeated Stackelberg security games: Learning with incomplete state information. *Reliability Engineering & System Safety*. 2020;195: 106695.
82. Alpern S, Morton A, Papadaki K. Patrolling games. *Operations research*. 2011;59(5):1246–57.
83. Garrec T. Continuous patrolling and hiding games. *European Journal of Operational Research*. 2019;277(1):42–51.
84. Klačka D, Kučera A, Lamser T, Řehák V. Automatic synthesis of efficient regular strategies in adversarial patrolling games. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*; 2018. p. 659–66.
85. Klačka D, Kučera A, Musil V, Řehák V. Regstar: efficient strategy synthesis for adversarial patrolling games. In: *Uncertainty in Artificial Intelligence*. PMLR; 2021. p. 471–81.
86. Portugal D, Iocchi L, Farinelli A. A ROS-based framework for simulation and benchmarking of multi-robot patrolling algorithms. In: *Robot Operating System (ROS)*; 2019. p. 3–28.
87. Portugal D, Rocha RP. Performance estimation and dimensioning of team size for multirobot patrol. *IEEE Intelligent Systems*. 2017;32(6):30–8.
88. Huang L, Zhou M, Hao K. Non-dominated immune-endocrine short feedback algorithm for multi-robot maritime patrolling. *IEEE Transactions on Intelligent Transportation Systems*. 2019;21(1):362–73.
89. Kunze L, Hawes N, Duckett T, Hanheide M, Krajník T. Artificial intelligence for long-term robot autonomy: A survey. *IEEE Robotics and Automation Letters*. 2018;3(4):4023–30.
90. Portugal D, Araújo A, Couceiro MSA, reliable localization architecture for mobile surveillance robots. In: *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2020;2020:374–9.
91. Chen T, Liu X, Xia B, Wang W, Lai Y. Unsupervised anomaly detection of industrial robots using sliding-window convolutional variational autoencoder. *IEEE Access*. 2020;8:47072–81.
92. Wardega K, Tron R, Li W. Resilience of multi-robot systems to physical masquerade attacks. In: *IEEE Security and Privacy Workshops (SPW)*. IEEE. 2019;2019:120–5.
93. Santana H, Ramalho G, Corruble V, Ratitch B. Multi-agent patrolling with reinforcement learning. In: *Autonomous Agents and Multiagent Systems, International Joint Conference on*. vol. 4. IEEE Computer Society; 2004. p. 1122–9.
94. Theile M, Bayerlein H, Nai R, Gesbert D, Caccamo M. UAV coverage path planning under varying power constraints using deep reinforcement learning. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE; 2020. p. 1444–9.
95. Luis SY, Reina DG, Marín SLT. A deep reinforcement learning approach for the patrolling problem of water resources through autonomous surface vehicles: The ypacarai lake case. *IEEE Access*. 2020;8:204076–93.
96. Huang Y, Zhu Q. Deceptive reinforcement learning under adversarial manipulations on cost signals. In: *International Conference on Decision and Game Theory for Security*. Springer; 2019. p. 217–37.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.