Contents

1 Introduction

Robots are increasingly leaving the confines of their highly structured and carefully curated environments within cages on manufacturing floors, academic laboratories, and purposefully arranged warehouses. This robot relocation is taking the robots to new places, where they are expected to operate across long temporal and spatial scales. For example, in precision agriculture, it is envisioned that robots will be persistently embedded in fields, tending to individual plants by monitoring and meeting their fertilizer, pesticide, or water needs [38, 381]. These agricultural robots will be present in the pastures throughout the full growing cycle, i.e., over an entire season [23]. Similarly, a number of environmental monitoring scenarios have been considered, where robotic sensor nodes are monitoring aspects of a natural environment [124, 392]. Examples include searching for the possibly extinct Ivory-billed Woodpecker in the forests of Louisiana [386], employing underwater robots for tracking marine pollution or the spread of invasive species [189, 407], or for monitoring the effects of climate change on the polar ice caps [388].

1.1 Long-Duration Autonomy

The deployment of robots over truly long time-scales in unstructured environments poses problems that are fundamentally different from those faced by robots deployed in factories or other controlled settings, where operating

Figure 1.1: Artist's portrayal of a NASA Mars Exploration Rover [196].

conditions exhibit only limited variability, power is readily available, and regularly scheduled maintenance routines ensure that minor technical problems do not accumulate to produce catastrophic failures. But, in *long-duration autonomy*, robots face a whole new set of challenges [71, 392], and this introductory chapter highlights some of the main themes and opportunities associated with these challenges, as well as makes the initial connection to ecology, i.e., to the *tight coupling between animal (robot) and its habitat (environment)*.

1.1.1 Lessons from Mars

When two Mars Exploration Rovers (MERs), MER-A and MER-B, landed on Mars in January 2004, they were tasked with completing individual missions spanning 90 Martian solar days, which corresponds to roughly 92.5 days on Earth [390]. Better known by their other names, *Spirit* and *Opportunity*, these rovers, as shown in Figure 1.1, managed to outlast their expected life spans by a significant margin and participate in five missions over 6 years and 2 months (*Spirit*) and a staggering 15 years and 1 month (*Opportunity*) [292].

Key to the longevity of these rovers was, of course, a great amount of highly ruggedized hardware and electronics, coupled with carefully designed, stress-tested, and clever engineering solutions. Additionally, the rovers had access to a virtually endless source of solar energy, and their solar arrays could generate as much as 140W per Martian day. Despite the abundance of energy, it was the power system that was expected to be the limiting factor in terms of the duration of the mission, as rechargeable batteries degrade over time and, as such, are no longer able to recharge to full capacity. But the real danger to the power system was the frequent Martian dust storms that not only would block the sunlight, but also accumulate dust on the solar panels, rendering them increasingly ineffective [296].

So why were *Spirit* and *Opportunity* able to perform their tasks significantly longer than expected? The answer was both simple and surprising. The same winds that sometimes caused dust storms on Mars would other times clean the solar panels by sweeping away dust [146]. These co-called "cleaning events" seem to have happened much more frequently than what NASA originally expected. As a result, the solar arrays were kept largely dust-free, and the life spans of the rovers were significantly extended—from less than a year to 15 years, in the case of *Opportunity*.

An immediate lesson one can draw from this interplanetary dust removal anecdote is that interactions between MERs and the environment proved to be beneficial to the rovers. But, at the same time, it was ultimately environmental factors that did the rovers in. *Spirit* got stuck in some particularly soft and sticky Martian soil during the summer of 2009. Despite efforts to free the rover, it was forced to reinvent itself as a stationary "science platform"—a task it performed for almost a year until contact was lost in 2010 [421]. *Opportunity*, on the other hand, did indeed get caught in a massive dust storm during the summer of 2018 that covered the solar panels so completely that it never recovered [421].

By necessity, the rovers were completely reliant on *in situ* solar energy, which, in turn, carried implications for how the robots functioned. One of the more striking manifestations of this dependence on sporadically present sunshine was how slowly the two MERs moved. *Opportunity*, which was the more peripatetic and well-traveled of the two rovers, had completed a full marathon on Mars by March 23, 2015, which translates to a rather leisurely finishing time of around 11 years and 2 months. The reason for this slow and steady pace can be traced back to considerations about energy conservation in conjunction with the need to stay away from trouble at all costs, as it was impossible to rescue a MER after a catastrophic event. As a result, the planning algorithms used for the rovers were highly conservative in terms of uncertainty management [74, 75, 257, 258]. Another way of phrasing this, using terminology borrowed from ecology, is that *survival* took precedent over most other considerations, including any notions of performance-based optimality.

The context in which this book is to be understood is that of long-duration autonomy, and the tale of the two impressive Mars rovers, *Spirit* and *Opportunity*, clearly highlights the two important themes of environmental interactions and survivability.

- Interactions between robot and the environment in which it is deployed play a key role in understanding design for long-duration autonomy; and
- Survivability, i.e., the explicit focus on avoiding getting caught in situations from which the robot cannot recover, takes precedent over all other design considerations.

It should be pointed out that although the MERs were absolute robotic marvels, and significantly advanced our understanding of robotics and autonomy, their operations were not what one would strictly call fully "autonomous." Instead, the rovers employed what NASA dubbed "directed autonomy," where commands were transmitted once per day to the rovers. The commands were encoded as event-driven sequences of motion commands that the rovers parsed using on-board stereo-vision and path-planning algorithms [50]. Despite this technicality, *Spirit* and *Opportunity* provide highly inspirational examples of robots that succeeded at carrying out a series of complex, long-duration missions over truly long time-scales.

1.1.2 Operations Beyond a Single Battery Charge

With the NASA Mars rovers as starting point, and using the key takeaways from their story, we have a handful of promising themes for characterizing and understanding long-duration autonomy. Perhaps the most important (and obvious) observation is that the robots have to be deployed over long periods of time for it to be considered "long-duration." One does not, however, need interplanetary travel to encounter situations where robots may be required to be deployed over long time-scales. In fact, our homes are increasingly being populated by household robots that are more or less in continuous operation, using dedicated charging or waste deposit stations. Environmental robots are being deployed in terrestrial or aquatic ecosystems to monitor factors such as plant growth, pollutants, wildfires, or climate trends, which may require the robots to be deployed for entire seasons. Warehouse robots are expected to perform fetch-and-carry operations; industrial robots are tasked with painting or welding; and mobile guide robots provide information to travelers in airports, art aficionados in museums, or patients in hospitals—all without taking breaks for maintenance or in other ways disrupting operations, e.g., [38, 187, 201, 381].

One way of defining long-duration autonomy is *deployment beyond a single battery charge* (or tank of gas), and where the recharging (or refueling) is part of the robot's portfolio of responsibilities.[1] Note that we phrased this in terms of "deployment" rather than in terms of a long-duration "mission." The reason for this is that we need to allow for situations where the mission may change, or where new missions may be requested. *Spirit* and *Opportunity* were sent to Mars to perform a focused science mission, but as they outlasted their expected life spans, they ended up performing in five different missions with completely different science objectives [292]. Perhaps even more striking and interesting is the situation where the robots may be deployed without any particular mission in mind at all. They are just asked to be present in an environment, waiting to be recruited to do whatever tasks need doing, following an *autonomy-on-demand* model, as opposed to a mission-centric view of what the deployment is supposed to be about [128, 304].

Regardless of whether the deployment involves a single, protracted mission, a sequence of multiple missions, or no clear mission at all,[2] two conditions must be satisfied for it to be considered a long-duration deployment, namely the deployment must last longer than a single battery charge, and the robot must be able to recharge itself.

- *Beyond a Single Battery Charge:* The scope of the deployment must be such that it is impossible for the robot(s) to successfully satisfy the requirements on a single battery charge; and

[1] We will use "battery" as shorthand for all sorts of different types of energy sources unless the context requires that the particulars be explicitly called out.

[2] One can of course argue—perhaps even successfully so—that having no mission at all is actually a mission in itself. As we will focus on "deployment" rather than "mission" as the defining characteristic of long-duration autonomy, this conundrum does not really matter for the developments in this book.

- *Autonomous Recharging:* No human intervention can be required in order for the energy sources to be replenished. Instead, the robot(s) must achieve this autonomously.

It is worth pointing out that the first condition, which states that a single battery charge is not sufficient, does not imply that clever power-management is not desired or needed.[3] On the contrary, power-management is certainly playing an exceedingly important role in the successful deployment of robots over long time-scales.

Once the robots are out in an environment for long periods of time, it is quite natural to draw inspiration from other "systems" that are present in environments over long time periods and need to "recharge," namely animals. This connection between animals and their habitats (ecology) and robots and their environments (henceforth known as "robot ecology") is indeed one of the central themes of this book. To this end, a number of biological organisms and habitats will be injected into the narrative in order to highlight and stress particularly salient ecological principles.

1.1.3 On the Value of Slowness

As already hinted at, the impetus behind the NASA Mars rovers' leisurely pace can be traced back to two primary reasons, namely the need to take it slow so as not to jeopardize the robots due to sudden or uncontrolled movements, and the need to conserve energy. As the saying goes, "slow and steady wins the race." Even though it is rare to actually see a tortoise and a hare line up and compete—if they did, the hare would most certainly win—the saying would indicate that the hare also runs a much higher risk of having something unforeseen happen to it due to its hasty outlook. Approaching new situations in hazardous, or even hostile, environments in a careful and deliberate manner is of particular importance when robots are supposed to be deployed over long time-scales, without human intervention. For instance, one of the primary reasons why underwater robotics is so tricky is that it is very hard and costly to recover malfunctioning or lost robots, e.g., [320, 367, 435]. Another manifestation of this idea can be found in the area of *safe learning*, which is

[3]Energy (joules) is what is available to the animal/robot, while power (joules per second, or watts) is the rate at which the energy is being delivered as work. So, "energy" will refer to the total charge of the battery, while "power" to the rate at which the battery is being drained.

predicated on the observation that a careless exploration of all state-action pairs can easily lead to the robot finding itself in disagreeable, and even harmful, configurations [5, 28, 45, 426].

Arguably, the primary reason for being slow—among animals as well as robots—is not to be cautious, but to conserve energy. As such, if the available energy is limited, which it usually is in nature, embracing a slow lifestyle can stretch the crucial energy resources further. For instance, *arboreal folivores* inhabit the ecological niche of spending their lives in the trees (arboreal), while sustaining themselves solely on leaves (folivore) [428]. This is a challenging strategy since in order to dwell productively among the trees, animals typically must be small and nimble so as not to simply fall down due to miscalculated leaps or broken branches. Now, contrast this arboreal size constraint with leaf-eating. Leaves are complicated foods in that they can be both toxic and structurally protected. In fact, as plants cannot move around in order to avoid their predators, they must come up with other means of defending themselves, like with thorns or spikes, or by chemical means [207]. Additionally, the cellulose fibers in the plant cell-walls that provide structural scaffolding to the leaves also make them hard to digest. As a result, animals who consume nothing but leaves must have a sufficiently long digestive tract, i.e., have a big enough gut, to break down these complicated foods [428]. The arboreal folivore is thus faced with the opposing requirements of being big enough to break down the food, yet small enough to live among the treetops.

What is the solution to this size dilemma faced by the arboreal folivores? Animals that occupy this ecological niche, such as koalas, two-toed and three-toed sloths, and some lemurs, all have roughly the same size, and they spend the vast majority of their time just sitting there among the treetops, doing nothing other than digesting their food. And when they do move, it is typically happening at an exceedingly slow pace. In other words, slowness has become a response to a severely energy-constrained existence. We will, throughout this book, return to these low-energy lifestyle animals as wellsprings of inspiration. In particular, the three-toed sloth will serve as a particularly suggestive source, culminating in Chapter 8 with the design of the SlothBot, a preview of which is shown in Figure 1.2.

For now, the takeaway from this initial discussion about power-management and slowness is simply that when operating in an environment where unlimited power is not available, and where the deployment specifications require the robot to function beyond a single battery charge, being slow

Figure 1.2: The SlothBot—a slow and energy-aware robot developed to perform environmental monitoring tasks—traverses a cable suspended between trees on Georgia Institute of Technology's campus.

is part of the toolbox. In fact, slowness is one of the design principles that separates long-duration autonomy from its short-duration counterpart.

1.2 Survivability

The shortest path between two points through a space populated by obstacles is obtained by moving as closely to the obstacles as possible [225, 237], as shown in Figure 1.3. Similarly, the fastest way for a car to come to a complete stop at a stop sign is maximal acceleration until the very last moment, and then the driver should slam on the brakes right at the stop sign [230].

Although optimal (minimum distance and minimum time, respectively), both of these strategies are problematic. What if the range-sensors used for detecting obstacles were not properly calibrated? In that case, the robot would hit rather than skirt the obstacles. Or, what if the model of the brake's effect on the car's motion was slightly wrong? In that case, the car might end up coming to a complete stop halfway through the intersection rather than at the stop sign, with potentially lethal consequences. As shown in [120, 208],

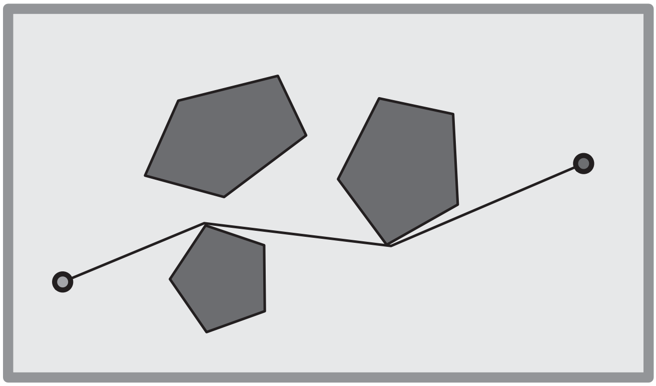

Figure 1.3: The shortest path through an environment typically runs as close as possible to obstacles (the polygons in the figure), thus rendering it non-robust to measurement errors [225, 236].

optimality and fragility are closely related concepts, meaning that optimal solutions are typically non-robust (or fragile) with respect to various types of disturbances, such as measurement or modeling errors. And, in long-duration settings, this lack of robustness can be catastrophic.

1.2.1 Costs and Constraints

Consider the problem of trying to make a solar-powered robot go to a particularly sunny spot to recharge the batteries. There are two different ways of achieving this objective. First, one can define a performance cost that evaluates how well the robot is progressing towards the goal location, e.g., by letting the cost be given by the distance to the goal. We denote this cost by $\mathscr{C}_{charge}$, and the controller should be chosen such that $\mathscr{C}_{charge}$ is minimized, either by incrementally moving against the gradient of the cost [251], or by finding the overall best strategy that minimizes the cost, e.g., by employing some path-planning method [225].

The second approach would instead be to define a constraint, abstractly encoded as $x \in \mathscr{G}_{charge}$, where x is the state of the robot, and where the constraint could say something like "the robot should always be able to make it back to the charging station given the current energy levels." The robot would then be allowed to move freely as long as it did not violate the constraint.

Additionally, the purpose of deploying this robot is likely more than simply making it go and recharge. Instead, the robot is probably expected to perform some primary task, such as detecting interesting events, protecting an area from intruders, or collecting Martian dust, and we let $\mathscr{C}_{task}$ be the cost that encodes how well (or poorly) this primary task is being performed. The question then arises: How should one balance these two different and potentially opposing requirements of recharging batteries and performing the primary task? The answer to this question depends on the context in which the overall mission is to be understood.

Consider, for example, the peculiar lifestyle exhibited by mayflies. Once they are past their nymph stage, they have an extremely short life span—as short as a few minutes for the female *Dolania americana*—and their single focus as adults is reproduction [186, 223]. This strategy leaves them with no need to feed and, as a result, they do not even have fully functioning mouthparts. Classifying this as long-duration autonomy would be a stretch. Similarly, if the robot, just like the mayflies, is supposed to execute the primary task and then be done, and the primary task can be completed with a single battery charge, then one should just ignore the charging requirement and solely minimize $\mathscr{C}_{task}$. One could call this approach the "short-duration autonomy" approach, and it is how robotics algorithms are typically approached. But, as already discussed, a key attribute of long-duration autonomy is that the robot should avoid catastrophic failures at nearly all costs since, once it has failed, there is no recovery. And, getting stuck somewhere in a dark corner of the deployment domain with completely depleted batteries certainly counts as a mission-ending failure. Borrowing, once again, from ecology, this translates to ensuring the survival of the robots.

With the notion of survival added to the mix, one approach could be to impose scheduled behaviors, e.g., periodic visits to the charging stations, during the execution of the primary task. The robot would thus switch between minimizing $\mathscr{C}_{task}$ and $\mathscr{C}_{charge}$ in response to a power-management scheme, e.g., [204, 374, 395]. But adhering to such fixed policies during execution without regard to implications on achieving task goals is not ideal as the robot would either solve the task it is supposed to solve, or recharge. But why not do both?

An attempt at doing both would be to encode survivability as a performance objective and somehow combine it with the other performance goals. This could, for example, be done via scalarization, where the overall performance cost would be given by a combination of the two costs, $\sigma\mathscr{C}_{task} + (1-\sigma)\mathscr{C}_{charge}$, for some $\sigma \in [0, 1]$. Alternatively, a so-called multi-objective optimization approach could be used. In the former case, the primacy of survivability is not ensured unless $\sigma = 0$, raising the possibility of catastrophic failures in the opportunistic pursuit of short-term gain. In the latter case, none of the available equilibrium or optimality concepts (e.g., Nash or Stackelberg equilibria, or Pareto optimality) ensure survival, instead balancing, in one way or another, the degree of survivability against other task performance criteria. In short, these types of approaches fail to adequately recognize (or exploit) the fact that *surviving is a prerequisite to thriving in long-duration autonomy applications*.

As will be seen in subsequent chapters, it is indeed possible to focus on the primary objective, yet ensure the survival of the robot. Constrained optimization provides the appropriate semantics for describing such an outcome. In other words, recognize the primary optimization or optimal control problem, but add in constraints that ensure survival by letting the robot solve,

$$\begin{aligned}&\text{minimize } \mathscr{C}_{task}\\&\text{subject to } x \in \mathscr{G}_{survive},\end{aligned} \tag{1.1}$$

where $\mathscr{G}_{survive}$ could be equal to $\mathscr{G}_{charge}$, or it could be a more general constraint that contains a number of other survival considerations as well, such as avoiding collisions or staying connected to other robots [128]. This seems like a highly promising way of abstractly capturing what long-duration autonomy could be about.

1.2.2 Robots that Do (Almost) Nothing

A particularly pertinent choice of performance cost, $\mathscr{C}_{task}$, in the previous section is to measure how much energy the robot is expending. In the absence of additional constraints, the optimal strategy would thus be to simply let the robot do nothing, i.e., to let the actuators exert no forces or torques on the system, which takes us close to the strategy employed by the arboreal folivores during long stretches of their existence. In fact, the conservation of energy is central to virtually all living organisms and, according to [357], the "purposeful expenditure of energy" is one key characteristic of what it means to be alive. As such, an initial, biologically motivated (yet mathematically vague and, for now, potentially ill-posed) attempt at formulating a design principle for robot ecology would be to modify the constrained optimization problem in Equation 1.1 to the following optimization problem,

$$\begin{aligned}&\text{do as little as possible}\\&\text{subject to } x \in \mathscr{G}_{survive}.\end{aligned} \tag{1.2}$$

One way of interpreting this formulation is as extreme, existential nihilism—the meaning of life is to expend as little energy as possible, while barely subsisting. Although a bit depressing, it is entirely consistent with basic ecological principles, where an animal's behavior is understood in large part through an energy-balance calculus [357].

Imagine now that a team of robots has been deployed in an environment, and are prepared to perform whatever tasks might be asked of them. One could call this setup *autonomy-on-demand*, and between tasks, the robots should merely be present in the environment, not running out of energy, and mostly do nothing. In such a scenario, the sloth is a good role model, and conducting oneself according to the constrained optimization problem in Equation 1.2 seems like a reasonable strategy. Once the robots are recruited to perform some task, the corresponding performance cost, $\mathscr{C}_{task}$, can be introduced, and the robot switches from the problem in Equation 1.2 to that in Equation 1.1.

But, beyond doing nothing, what the autonomy-on-demand framework suggests is the possibility of having robots with free time. The question then becomes, what should these robots spend their time doing? There are indeed opportunities afforded by being a robot of leisure. It could, for example, improve its skills by learning and exploring better control policies. It could also learn completely new skills. This is a bit more delicate as the robot is literally tasked with doing nothing, and most of the machine learning apparatus requires some sort of goal or reward against which the suitability of the control actions can be evaluated. In the absence of such goals, one instead needs to move towards a more "curiosity-driven" learning paradigm [217, 393, 416], where the robot explores state-action pairs without a predefined, clear goal, or where mismatches between actual and modeled effects are being pursued for the purpose of getting more accurate models of the robot's capabilities. But, perhaps most importantly, interactions between environment and robot can be better understood.

The interactions between robot and its habitat is imperative to the robot ecology framework in that this coupling must be understood and leveraged in order for the robot to successfully dwell in an environment over sustained periods of time. This is the topic of the next section, and an example is shown in Figure 1.4, where a robot, using computer vision, must learn to discriminate between objects according to their texture and color profiles. For instance, the robot should learn to tell tall grass from boulders, as its ability to traverse these "objects" is completely different. And, the only way to gauge the

Figure 1.4: Based on an object's texture and color properties, a robot must learn which objects in the environment can be traversed and which cannot.

"traversability" of a particular type of object is to interact with it, e.g., to try to drive through it without getting stuck [396].

1.3 Coupling Between Environment and Robot

As already discussed, survivability, i.e., the ability to avoid situations where survival can no longer be ensured, can be naturally encoded as a constraint rather than as a performance objective. This way of formulating survivability is also consistent with ecological principles, where richness of behavior is a direct function of environmental constraints [308, 357, 385], including the abundance and distribution of resources, favorable microclimates, and the prevalence of suitable mates or predators. Indeed, when ecologists study the distribution of species and the composition of populations and communities, the environmental reality and the associated ecological constraints are as important, if not more so, than any "goal-driven" behaviors [357, 385].

Based on this observation that constraints are fundamentally important to animal behavior, one can thus ask if the constraint-based vantage point translates to effective control design principles for engineered systems as well. As such, we will approach the design problem as one where the robots' behaviors are mostly constraint-driven, such as avoiding collisions with obstacles or other robots, or never completely depleting the batteries, as opposed to goal-driven. In fact, these types of constraints can be derived (albeit subject to a slight robotic reinterpretation) from basic ecological principles. As ecology is aimed at understanding the interaction of organisms with their environments and with other organisms, this is a particularly fruitful metaphor also for robots leaving the highly curated laboratory or factory settings, and entering dynamic, unstructured, natural environments across long temporal and spatial scales.

1.3.1 Ecosystems

As was discovered by the Mars rover team, the connection between robot and environment was even more important for the longevity of the robots than what was originally thought. Not only was the environment a source of energy, it was an existential threat through dust build-up and soft and sticky sand. But it also provided unexpected help when the Martian winds would swipe the solar panels clean, thereby overcoming other, more adverse environmental factors. What this anecdote tells us is that the robot and the environment it inhabits should be thought of as a single system, which brings us within striking distance of the idea of an ecosystem.

In the 1930s, a vibrant discussion took place among ecologists about the proper way of thinking about these interconnections, and the term "ecosystem" was coined by the British ecologist A. G. Tansley in 1935. He writes [402]: "The more fundamental conception is, as it seems to me, the whole system (in the sense of physics), including not only the organism-complex, but also the whole complex of physical factors forming what we call the environment of the biome—the habitat factors in the widest sense. *Though* the organisms may claim our primary interest, when we are trying to think fundamentally we cannot separate them from their special environment, with which they form one physical system." This way of thinking about organism (animal/robot) and environment as part of the same system—not in a loose, metaphorical sense but in a tight, physical sense—will prove to be a fruitful way of approaching long-duration deployments. In fact, as animals and plants live (literally) in a physical environment, their *form* and *function* must obey the rules of the physical world [385].

An illustrative example of how form is determined by the physical environment is the size of the pores in avian eggshells. As gas is constantly passing through the eggshell throughout the incubation period to deliver oxygen and nutrients, the movement of the gas follows a diffusion process, which means that environmental factors such as altitude, temperature, and humidity all matter to the type of egg (size of pores) the bird lays [357]. Similarly, rates of processes (r) and animal dimensions (d) typically satisfy an "allometric" relationship, $r = ad^b$, with b being the *allometric constant* [86]. For instance, heart rate versus body mass has an allometric constant of $b \approx -0.2$ among mammals, while the metabolic rate vs. body mass has $b \approx 0.7$ [86, 357]. In other words, chipmunks have a higher heart rate but a lower metabolic rate than elephants. And these environmentally informed form factors have implications for the animals' functions.

On the functional side, animals constantly move among so-called environment patches as the environment changes, over days, months, and even years [357, 385]. And, for the purpose of this book, functional considerations will play a more prominent role than form considerations. That is not to say that form does not matter—it does. Only that the focus of this book is on the control design considerations when deploying robots over long time-scales, i.e., hardware will play second fiddle to software.

There are a number of situations where this idea of functional coupling between robot and environment is not only useful, but crucial when deploying robots over long time-scales. The most apparent and covered situation is the recharging of batteries using energy from the sun, meaning that the robot must, every so often, find itself in a place with ample sunlight. In other words, cave-dwelling robots must either surface every now and then to bask in the sun, or they must rely on some other source of energy. However, energy harvesting does not provide the only beneficial coupling between robot and environment. When aerial gliders or marine robots move through their domains, updrafts and ocean currents, respectively, provide opportunistic sources of low-cost mobility [6, 147, 373]. And passive walkers, e.g., [432], only really function in worlds consisting solely of gentle downhill slopes—in all directions.

1.3.2 Natural and Engineered Environments

So far, the discussion has been focused on natural environments. But, the importance of harnessing the coupling between robot and its habitat is certainly not diminished in engineered environments. Robots roaming around in warehouses or homes must not only be able to locate outlets, they must be equipped with the proper hardware (e.g., plugs) to allow them to take advantage of the available energy sources.

From a mobility perspective, environment and hardware design go hand-in-hand, and one reason why humanoid robots are deemed particularly useful as companion robots in our homes is that our domestic environments are already built for bipedal humanoids that are four- to six-feet tall, with staircases that reward legged locomotion and doorknobs that are strategically placed at certain heights. In other words, our homes are already superbly well-suited for humanoid robots [156]. One does not, however, have to look to the world of autonomous robots to see this phenomenon. For more familiar and mundane vehicles, a carefully engineered environment is oftentimes called for. Trains are extraordinarily well-positioned to take advantage of train tracks, airplanes of airport runways, and cars of highways. We even sometimes modify animals to make them fit our engineered environments, such as putting shoes on horses or electric collars on dogs.

As we will see in a later chapter, this idea of making slight modifications to the environment for the explicit purpose of rendering the robots' existences more productive and safe will prove beneficial when deploying robots up in the treetops in persistent environmental monitoring applications. Climbing, as practiced by arboreal animals, is problematic from a safety point of view in that a robot that falls out of a tree will probably not be able to continue on with its mission. But, by stringing cables in the treetops, we can ensure that the robots can dwell successfully in the tree canopies and remain safely suspended, even when actuators fail or energy levels drop precipitously.

Regardless of whether the robots are to be deployed in jungles, on train-tracks, or in kitchens, when the deployments transpire over long temporal scales, unexpected things are inevitable [71]. The world is fundamentally a messy place, and any attempt at enumerating all the possible things a robot might encounter, in all but the most sterile environments, is doomed to fail. As such, the strategies employed in long-duration settings must support adaptation to new situations. And they must achieve this while ensuring the robot's safety at all times and at (almost) all costs.

1.4 Summarizing and Looking Ahead

What this introductory chapter has done is paint a mood picture and describe some of the challenges associated with long-duration autonomy. It also identified a collection of guiding principles that permeate the book and that illustrate why long-duration autonomy, as compared to its short-duration counterpart, is different from a control design vantage point.

These principles combine together under the umbrella of *Robot Ecology* as follows:

- The tight coupling between robot and environment is not only important, it is absolutely crucial if the robots are to exhibit longevity;
- Survival is a prerequisite to thriving, i.e., rather than minimizing performance-based costs, the control design should focus on ensuring that various safety constraints are satisfied;
- Key among safety constraints is power-management since, if the robot finds itself with depleted batteries without any ability to recharge, it is game over;
- When the deployment takes place over truly long time periods, being fast is oftentimes both energetically wasteful and dangerous, and the robots should embrace a slow lifestyle; and
- Any attempt at enumerating everything the robot may encounter is doomed to fail, and the ability to adapt to changing environmental conditions and missions is a core attribute in long-duration autonomy.

The remainder of this book will take these casual observations and make them more precise and mathematically well-defined. Part II of the book will serve this purpose by establishing control barrier functions (CBFs) as the proper framework for talking about robot survival using the constrained optimization semantics. Once CBFs have been adequately introduced, they will be drawn on to support persistified robot tasks, i.e., to modify nominal controllers in order to extend the robots' life spans indefinitely. The developed tools and techniques will then be employed in Part III for long-duration deployments in a number of different settings, with a particular focus on environmental monitoring and conservation tasks. However, before these tools can be unleashed, a more thorough discussion is needed of what robot survival actually entails, and of how ecological principles can be put to use towards the overarching theme of establishing a theory of robot ecology.

Index