

Energy-Based Models for Virtual Creatures

Junior Rojas

jrojasdavalos@gmail.com

Abstract

This paper presents an energy-based approach for simulating virtual creatures, advocating for a shift from traditional monolithic physics engines to a more flexible implementation approach centered on energy minimization and automatic differentiation. By integrating insights from established disciplines alongside emerging concepts such as scale-free cognition, this approach enables a comprehensive modeling of behaviors, where everything from basic physical phenomena, such as inertia and elasticity, to more complex behaviors, such as robust locomotion, can be interpreted as goal-directed behavior.

Introduction

In simulations of virtual creatures, capturing complex, life-like behaviors is significantly enhanced by the ability to model a wide range of physical phenomena and material properties (Cheney et al., 2014; Kriegman, 2019). While physics engines have traditionally been crucial to these simulations, they tend to prioritize the performance and accuracy of established physical models, which often results in monolithic systems that constrain the potential for exploring more diverse physical behaviors.

Simulations of virtual creatures, with their inherent need for flexibility and adaptability, stand to benefit significantly from a departure from traditional physics engines. Instead of viewing them as rigid enforcers of physical laws, they can be reconceptualized as collections of energy functions that can be easily composed to represent a wide range of behaviors. This reconceptualization facilitates a more fluid integration of cognitive processes, potentially leading to a unified framework where cognition and physics are seamlessly intertwined.

The term “energy” is prevalent in various fields, including physics, machine learning, and biology (LeCun et al., 2006; Friston et al., 2006; Friston, 2012). In this paper, the term “energy” is initially used to refer to potential energy specifically, and then its definition is broadened to refer to any scalar-valued function that measures compatibility between variables as a means of implicitly capturing their dependencies (LeCun et al., 2006). The following sections elaborate




Figure 1: Entities with large cognitive light cones can leverage the inherent goals of simpler entities, such as the tendency of springs to maintain their rest lengths. For example, by dynamically adjusting setpoints in the form of spring rest lengths, it is possible to dynamically reshape the energy function to achieve robust locomotion. Despite the introduction of a neural network to adjust the setpoints, the simulation remains primarily driven by energy minimization, where the energy function is composed of six differentiable terms: muscles, triangles, gravity, collision, friction, and inertia. Implementation available at <https://github.com/juniorrojas/algovivo>.

on the specific interpretation of “energy” within the context of virtual creature simulations. A concrete implementation (Rojas, 2023) is used to further clarify and exemplify the concepts discussed in this paper.

Building on emerging concepts such as cognitive light cones and scale-free cognition (Levin, 2019, 2022; Biswas et al., 2021; Zhang et al., 2023), this paper emphasizes the spatio-temporal boundaries within which cognitive individuals operate. These concepts establish that each cognitive individual is defined by a computational boundary, a delimited region of space-time within which they can perceive, model, and influence their surroundings. Furthermore, they challenge conventional distinctions between cognitive and physical processes, suggesting a continuum where goal-directedness extends from the simplest physical phenomena to complex cognitive behaviors.

Energy Minimization

The approach advocated in this paper is grounded in the principle of energy minimization, inspired by the intuitive clarity and simplicity that scalar loss functions provide in neural network models. Much like the loss function encapsulates the objective of a neural network in a single scalar value despite its high-dimensional nature, potential energy functions offer a scalar representation of the objective of a physical system.

This approach simplifies the comprehension of physical systems as goal-directed entities naturally evolving towards states of minimum energy and also addresses the challenge of deciphering high-dimensional force vectors akin to neural network gradients. The force can be straightforwardly derived as the negative gradient of the potential energy. By adopting the energy function as the primary definition, this approach can benefit from recent advances in automatic differentiation (Moses and Churavy, 2020; Paszke et al., 2019; Bradbury et al., 2018) to circumvent the complexity of explicit force models, offering a clearer framework for understanding system dynamics.

Dynamics as energy minimization

The concept of energy minimization, foundational to understanding static configurations and dissipative processes, can be extended to encompass the dynamics of motion through the inclusion of an inertia-related “energy” term. This addition represents one key conceptual shift: dynamics, conventionally expressed through forces and accelerations and formalized with second-order differential equations, can be reformulated as a process of minimization. Specifically, dynamics can be incorporated by introducing a penalty term that accounts for deviations from inertial motion as a way to quantify how much a system’s trajectory diverges from what would be expected based purely on its initial momentum. This approach is rooted in the methodologies of implicit numerical integration methods such as backward Euler, renowned for their stability and efficacy in physics-based simulation (Baraff and Witkin, 1998; Gast et al., 2015).

While the term inertial “energy” is used in this context, it is crucial to distinguish it from kinetic energy. This usage aligns with the terminology in energy-based models (LeCun et al., 2006), where “energy” more broadly refers to a scalar-valued function that implicitly captures dependencies between variables (here, the dependency between successive states in a time step).

The addition of this inertia penalty term captures the essence of Newton’s first law of motion within the framework of energy minimization. Each vertex, a basic unit of simulation, can be thought of as having a simple temporal light cone, which represents its ability to “remember” its momentum and attempt to follow its natural inertial path. This memory, while basic, provides each vertex with a rudimentary form of goal-directedness focused on maintaining

its momentum.

By incorporating additional energy functions into the simulation, vertices are required to find a balance between following their natural inertial paths and seeking states of lower potential energy, such as those influenced by gravity. Moreover, this approach reconciles conservative and non-conservative forces, such as friction, framing system dynamics as an ongoing process aimed at minimizing a function composed of multiple energy terms.

Exploiting energy minimization to achieve large-scale goals

In this energy-based approach, multi-vertex potentials also exemplify how simple systems, driven by energy minimization, can exhibit spatial awareness and intentionality. Springs that maintain their rest length and elastic triangles that preserve their shape and area demonstrate very basic forms of goal-directed behavior. These systems, governed by the need to maintain certain energetic configurations, can be associated with rudimentary cognitive light cones, where an entity’s awareness extends only as far as the vertices it encompasses.

Building on this foundation, entities with larger cognitive light cones can exploit the intrinsic goals of simpler entities, such as a spring’s inherent drive to maintain its rest length, to pursue larger-scale goals. By introducing mechanisms that allow for the adjustment of these setpoints (effectively, the goals of simpler entities) entities with larger cognitive light cones can exert nuanced influence over their surroundings. The model illustrated in Figure 1 demonstrates the capability of an agent to manipulate springs that act like muscles by dynamically reshaping the energy function to achieve robust locomotion over long time horizons. Despite the introduction of a neural network to adjust the setpoints, the simulation remains primarily driven by energy minimization, where the energy function is composed of six differentiable terms: triangles, muscles, gravity, collision, friction, and inertia.

Towards practical implementations

This paper deliberately avoids an in-depth discussion of the neural network training process, which traditionally receives ample attention as a goal-directed process. Instead, it focuses on illustrating how energy minimization principles can redefine basic physical phenomena as goal-directed behaviors. This approach facilitates practical implementations that effectively utilize modern automatic differentiation tools and leaves room for further development of new energy functions to more comprehensively model a broader spectrum of behaviors. For a concrete example of this energy-based approach, the reader is referred to (Rojas, 2023), which implements the six energy functions used in the simulation illustrated in Figure 1.

References

- Baraff, D. and Witkin, A. (1998). Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, page 43–54, New York, NY, USA. Association for Computing Machinery.
- Biswas, S., Manicka, S., Hoel, E., and Levin, M. (2021). Gene regulatory networks exhibit several kinds of memory: Quantification of memory in biological and random transcriptional networks. *iScience*, 24(3).
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs. <https://github.com/google/jax>.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2014). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. *SIGEVolution*, 7(1):11–23.
- Friston, K. (2012). A free energy principle for biological systems. *Entropy*, 14(11):2100–2121.
- Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1-3):70–87.
- Gast, T. F., Schroeder, C. A., Stomakhin, A., Jiang, C., and Teran, J. M. (2015). Optimization integrator for large time steps. *IEEE Trans. Vis. Comput. Graph.*, 21(10):1103–1115.
- Kriegman, S. (2019). Why virtual creatures matter. *Nature Machine Intelligence*, 1(10):492–492.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. J. (2006). A tutorial on energy-based learning. *Predicting Structured Data*, 1.
- Levin, M. (2019). The computational boundary of a “self”: Developmental bioelectricity drives multicellularity and scale-free cognition. *Frontiers in Psychology*, 10.
- Levin, M. (2022). Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16.
- Moses, W. and Churavy, V. (2020). Instead of rewriting foreign code for machine learning, automatically synthesize fast gradients. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12472–12485. Curran Associates, Inc.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA.
- Rojas, J. (2023). Algovivo: An energy-based formulation for soft-bodied virtual creatures. <https://github.com/juniorrojas/algovivo>.
- Zhang, T., Goldstein, A., and Levin, M. (2023). Classical sorting algorithms as a model of morphogenesis: self-sorting arrays reveal unexpected competencies in a minimal model of basal intelligence.