

Computational Food Materials Design of Plant Polysaccharide-Based Cryoprotectants for Frozen Bakery Applications

Saumya Cheema

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
scheema@missouri.edu

Marcelina Imiolczyk

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV,
USA.
imiolczyk1986@unr.edu

Aleks Dylak

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
dylak1996@colostate.edu

Abstract

Frozen bakery systems occupy a critical position in contemporary food supply chains because they reconcile centralized production with distributed retail baking, reduce waste, and support consistent product quality. Their performance, however, remains constrained by ice formation, recrystallization, dough weakening, yeast inactivation, and structural collapse during freeze-thaw cycles. Plant polysaccharide-based cryoprotectants offer a promising alternative to conventional hydrocolloids and synthetic additives, yet their rational design remains difficult because cryoprotection emerges from coupled molecular, colloidal, rheological, and process-scale phenomena. This paper examines computational food materials design as a system-level architecture for discovering and deploying plant polysaccharide cryoprotectants in frozen bakery applications. It argues that the field must move beyond isolated molecular screening toward integrated platforms that connect polymer structure, hydration behavior, sol-gel transitions, dough rheology, fermentation performance, freeze-thaw robustness, and baking outcomes. The discussion emphasizes structural trade-offs among ice inhibition, dough extensibility, gas retention, cost, sensory neutrality, and regulatory acceptance. It also considers data infrastructure, digital twins, edge and cloud deployment, model governance, fairness across suppliers and regions, and sustainability implications. A central claim is that computational design becomes industrially meaningful only when it is coupled to traceable data, interpretable models, robust process control, and inclusive governance. The paper concludes that plant polysaccharide cryoprotectants should be treated not merely as additives but as components of socio-technical infrastructure whose value depends on reproducibility, scalability, policy alignment, and equitable access.

Keywords

computational food design; plant polysaccharides; cryoprotectants; frozen dough; bakery systems; digital twins; machine learning; food governance; sustainability.

1. Introduction

Frozen bakery products have become central to modern food systems because they allow dough, par-baked goods, and laminated pastries to be produced at scale, transported through cold chains, and finished close to the point of consumption. This architecture reduces skill requirements at retail sites, extends shelf life, and enables firms to respond to regional demand without maintaining full production capacity everywhere. The technical foundation of this system is cryopreservation of dough structure, yet freezing remains intrinsically damaging. Ice crystals disrupt gluten networks, reduce yeast viability, alter water distribution, and promote recrystallization during storage and temperature fluctuation. These effects have been documented across frozen dough research, where storage time, freeze-thaw cycling, and formulation choices influence rheology and baking quality [1], [3], [4]. The result is a persistent trade-off between operational convenience and product quality.

Cryoprotectants are used to mitigate these effects by modifying water activity, glass formation, ice crystal habit, and biopolymer interactions. Conventional approaches have relied on sugars, polyols, hydrocolloids, emulsifiers, and ice-structuring proteins, but each class carries limitations. Sugars can alter sweetness and fermentation, hydrocolloids can produce undesirable gumminess, and proteins may introduce allergenic or cost constraints. Plant polysaccharides are attractive because they are abundant, often biodegradable, and chemically diverse. They can be sourced from agricultural side streams, including pectin-rich peels, cereal brans, seed gums, and hemicellulosic fractions. Their cryoprotective potential arises from hydration, viscosity, gelation, and colloidal stabilization rather than from a single molecular mechanism. This complexity makes them difficult to optimize empirically but also makes them suitable for computational exploration. The broader scientific context of glassy-state stabilization and water dynamics provides a foundation for understanding why some polysaccharides protect frozen systems while others fail [1], [2], [6].

Computational food materials design promises to accelerate discovery by integrating molecular simulation, machine learning, process modeling, and experimental feedback. In materials science, data-driven approaches have already transformed molecular and materials discovery, inverse design, and property prediction [10], [11], [12], [13]. Food systems, however, introduce additional layers of complexity. A cryoprotectant must not only inhibit ice growth but also remain compatible with dough handling, fermentation, baking, sensory expectations, labeling rules, and industrial cost structures. Therefore, the design problem is not simply one of molecular property optimization. It is a system architecture problem in which computational models, data pipelines, physical experiments, manufacturing assets, and governance regimes must be aligned. Recent work on plant polysaccharide fractions in frozen dough demonstrates that cryoprotection can be linked to sol-gel transition behavior and structural organization, suggesting that rational design is possible when molecular structure and process context are considered together [9].

This paper develops a system-level perspective on computational food materials design for plant polysaccharide-based cryoprotectants. It first frames cryoprotection as a coupled materials and infrastructure problem. It then examines plant polysaccharides as tunable architectures and discusses the trade-offs that govern their use. Subsequent sections address computational workflows, data infrastructure, digital twins, deployment architectures, robustness, fairness, governance, sustainability, and policy. The central argument is that computational design must be embedded in reproducible data ecosystems and accountable deployment pathways if it is to produce cryoprotectants that are not only effective in the laboratory but also viable in distributed frozen bakery operations.

2. Cryoprotection as a Materials and Systems Problem

Cryoprotection in frozen bakery systems cannot be reduced to a single physical variable. Freezing concentrates solutes, changes ionic strength, dehydrates biopolymers, and creates mechanical stresses as ice expands. During frozen storage, temperature fluctuations drive recrystallization, which enlarges ice crystals and increases damage to gluten and starch structures. The glassy state can slow molecular mobility and improve stability, but glass formation must be balanced against dough workability and fermentation requirements [1], [6]. Water activity and phase transitions thus frame the problem, but they do not fully determine baking performance [2]. A formulation that protects yeast may impair gas retention, while a formulation that strengthens dough may reduce extensibility.

From a systems perspective, cryoprotectant design must satisfy multiple interfaces. At the molecular interface, polysaccharides interact with water, ions, proteins, and starch through hydrogen bonding, electrostatic association, and excluded-volume effects. At the colloidal interface, they influence protein aggregation, starch granule integrity, and emulsion stability. At the rheological interface, they alter viscoelasticity, extensibility, and proofing behavior. At the process interface, they must survive mixing, sheeting, freezing, storage, thawing, and baking. At the economic interface, they must be cost-effective at industrial dosage and compatible with existing cold-chain assets. These interfaces are not independent. A change in molecular weight may improve ice inhibition but increase viscosity beyond process limits. A change in charge density may stabilize proteins but create undesirable interactions with salts or enzymes.

The literature on polyelectrolyte complexation and coacervation illustrates how colloidal assembly can be tuned, but it also shows that small changes in pH, ionic strength, or polymer composition can shift behavior dramatically [8]. In frozen dough, this sensitivity is amplified because freezing itself changes the solvent environment. A computational model that ignores such coupling may predict a promising candidate that fails under realistic freeze-thaw cycles. Similarly, empirical optimization may identify a working formulation without revealing why it works or how to transfer it to another flour, water source, or bakery process. The required approach is therefore hybrid: mechanistic models provide interpretability, data-driven models capture nonlinear interactions, and experimental platforms provide grounding. This aligns with broader trends in computational chemistry and materials informatics, where predictive insight depends on combining simulation, learning, and domain knowledge [12], [13].

The systems challenge also includes biological variability. Wheat cultivars, flour aging, yeast strains, enzyme activity, and water mineral content differ across regions and seasons. A cryoprotectant that performs well in one supply chain may fail in another. Computational design must therefore treat robustness as a first-class objective, not an afterthought. Robustness means acceptable performance across a distribution of inputs, not optimal performance at a single laboratory condition. It also means graceful degradation when storage temperatures deviate, when fermentation times vary, or when ingredient substitutions occur. This framing connects cryoprotectant design to resilience engineering and infrastructure planning, because frozen bakery quality is ultimately produced by a distributed socio-technical system.

3. Plant Polysaccharides as Tunable Cryoprotective Architectures

Plant polysaccharides are structurally diverse and functionally plastic. Pectins, arabinoxylans, beta-glucans, galactomannans, xyloglucans, and hemicellulosic fractions differ in sugar

composition, branching, charge, molecular weight, and chain conformation. These features determine their hydration behavior, viscosity, gelation, and ability to interact with proteins and starch. Some polysaccharides form cohesive networks that stabilize dough structure, while others act as water-binding agents that reduce ice crystal growth. Still others may interfere with gluten development or yeast activity if used at inappropriate concentrations. The design space is therefore broad but constrained by food compatibility, sensory neutrality, and regulatory status.

The cryoprotective mechanism of plant polysaccharides is often described through water-binding and glass-forming effects, but emerging evidence points to more specific structural roles. Sol-gel transitions can create transient networks that organize water and limit ice propagation without making dough unprocessable [9]. Such transitions may be sensitive to polymer concentration, temperature history, ionic environment, and the presence of other biopolymers. A polysaccharide fraction that forms a protective gel during freezing may need to relax during proofing to allow gas expansion. This dynamic requirement suggests that cryoprotectants should be designed as responsive materials rather than static additives. The goal is not maximum gel strength but timed and reversible structuring.

Structural trade-offs are central. High molecular weight and extensive branching can increase viscosity and water holding, which may reduce ice damage, but they can also inhibit gluten network formation and reduce loaf volume. Charged polysaccharides may stabilize proteins through electrostatic interactions but can also promote aggregation or phase separation depending on pH and salt. Neutral polysaccharides may be more compatible with dough but less effective at controlling ice. Mixtures can outperform single polymers by combining complementary mechanisms, but mixtures also complicate characterization, quality control, and regulatory review. Computational design can help navigate these trade-offs by predicting structure-function relationships and identifying Pareto-optimal formulations across competing objectives.

Plant sourcing introduces another layer of complexity. Agricultural side streams are attractive for sustainability, but their composition varies with crop, season, extraction method, and geography. A computational platform that assumes fixed input properties will produce models that fail in real supply chains. Instead, the platform must represent raw material variability as part of the design space. This can be achieved through probabilistic models, robust optimization, and continuous data feedback from extraction and quality control. The resulting formulations may be less elegant than single-molecule solutions, but they may be more deployable because they tolerate real-world variation. This perspective shifts attention from isolated molecules to material ecosystems.

4. Computational Design Workflow and Data Infrastructure

A credible computational workflow begins with a clear representation of the design problem. Inputs include polysaccharide structure, composition, molecular weight distribution, charge, branching, source, and processing history. Outputs include cryoprotective efficacy, dough rheology, fermentation performance, freeze-thaw stability, baking quality, sensory attributes, cost, and regulatory compatibility. Between inputs and outputs lie multiple models: quantum and molecular simulations for local interactions, coarse-grained models for polymer dynamics, continuum models for heat and mass transfer, and machine learning models for property prediction. No single model can span all scales, so integration is essential. Machine learning for molecular and materials science has shown that predictive models can accelerate discovery when they are trained on high-quality, diverse data [10], [12]. Inverse design and

generative models further expand the space by proposing candidates that meet target properties [11]. In food systems, however, targets are multi-objective and context-dependent.

Data infrastructure is therefore as important as model architecture. A useful platform must manage experimental protocols, raw measurements, metadata, provenance, and versioning. It must distinguish between intrinsic material properties and process-dependent outcomes. It must support reproducibility across laboratories and pilot plants. Without such infrastructure, computational models become isolated demonstrations rather than cumulative scientific assets. The field also needs standard vocabularies for cryoprotection, dough handling, and baking quality so that datasets from different sources can be integrated. This is analogous to materials informatics, where shared schemas and benchmark datasets have accelerated progress [12], [13]. In food, the challenge is greater because sensory and regulatory dimensions resist simple numerical encoding.

A practical workflow can be organized as a closed loop. Computational models generate candidate formulations or structural modifications. These candidates are tested through high-throughput or medium-throughput experiments. Results are curated into structured datasets. Models are retrained or recalibrated. Promising candidates advance to pilot-scale freezing, storage, and baking trials. Industrial feedback informs the next iteration. This loop must be governed by uncertainty quantification, because overconfidence in sparse data can lead to costly scale-up failures. Interpretability is also critical. Bakers and process engineers need to understand why a formulation is expected to work, especially when it changes handling or fermentation behavior. Interpretable models can support trust, troubleshooting, and regulatory dialogue.

The computational design of plant polysaccharide cryoprotectants also requires attention to negative results. Failed formulations are informative because they reveal boundaries of the design space. If only successful experiments are recorded, models develop survivorship bias and overestimate performance. Data infrastructure should therefore capture failures, outliers, and inconsistent batches. This is a cultural as well as technical challenge. Academic incentives often reward positive findings, while industrial secrecy limits data sharing. Hybrid governance models, such as precompetitive consortia with shared benchmarks and protected proprietary formulations, may be necessary to balance openness and commercial interest.

5. Digital Twins, Process Control, and Deployment Architecture

Once a cryoprotectant is formulated, its performance depends on deployment. Frozen bakery production is distributed across mixing, dividing, shaping, freezing, storage, transport, thawing, proofing, and baking. Each stage has its own time-temperature profile and mechanical stresses. A digital twin can represent these stages as a connected system, allowing virtual experimentation and real-time monitoring [14], [15]. In this context, a digital twin is not merely a simulation. It is a continuously updated model of a physical process, informed by sensors, historical data, and human observations. It can predict how a dough batch will respond to a delayed freezer entry or an unexpected temperature spike. It can also recommend adjustments to proofing time or baking temperature.

Deployment architecture determines whether such capabilities are practical. Cloud platforms offer scalable computation and centralized model updates, but they introduce latency, connectivity dependence, and data governance concerns. Edge computing allows real-time control at the bakery or cold-chain node, but it constrains model size and complicates maintenance. A hybrid architecture is often appropriate: latency-critical control runs locally,

while fleet-level learning and model retraining occur in the cloud. This creates trade-offs among cost, privacy, robustness, and update speed. For small bakeries, the cost of sensors and digital infrastructure may be prohibitive, raising fairness concerns. If computational cryoprotectants are optimized only for large, instrumented facilities, they may deepen inequalities in the baking sector.

Process control must also account for human practice. Bakers often rely on tactile and visual cues that are not easily captured by sensors. A digital twin that ignores tacit knowledge may produce recommendations that are technically optimal but operationally implausible. Effective deployment therefore requires participatory design, where bakers, millers, ingredient suppliers, and equipment manufacturers help define data streams and user interfaces. The same cryoprotectant may require different usage guidance depending on mixer type, freezer design, and local water chemistry. Computational models can support adaptation, but only if they are embedded in workflows that respect local constraints.

The digital twin concept also extends to the cold chain. Temperature loggers, telematics, and smart packaging can provide data on storage conditions from factory to retail. If these data are integrated with dough formulation models, it becomes possible to predict quality loss and dynamically adjust distribution. This could reduce waste and improve consistency. However, it also raises questions about data ownership and liability. If a quality failure is predicted, who is responsible for corrective action? If data are proprietary, how can risks be shared across the chain? These are governance questions as much as technical ones. Deployment architecture must therefore include clear rules for data access, model accountability, and dispute resolution.

6. Robustness, Fairness, and Governance

Robustness in computational cryoprotectant design means maintaining performance under uncertainty. Sources of uncertainty include raw material variability, measurement error, model error, process drift, and unexpected events such as power outages or cold-chain disruptions. A robust formulation performs acceptably across a range of conditions rather than optimally in a narrow window. This requires stress testing, sensitivity analysis, and scenario planning. It also requires models that can detect when they are outside their training domain and default to conservative recommendations. In frozen bakery systems, robustness has direct economic value because it reduces recalls, waste, and customer complaints.

Fairness enters through access, representation, and distribution of benefits. Plant polysaccharide sources are often agricultural by-products, and their valorization could benefit rural economies. However, if extraction and computational infrastructure are concentrated in a few firms or regions, value may accrue disproportionately to powerful actors. Small bakeries may lack the capital to adopt digital twins or specialized formulations. Consumers may face higher prices or fewer choices if proprietary cryoprotectants dominate the market. Fairness also concerns data. If models are trained primarily on wheat varieties, processes, and consumer preferences from high-income regions, they may perform poorly elsewhere. Inclusive datasets and participatory governance can mitigate these risks.

Governance frameworks for algorithmic systems emphasize transparency, accountability, and contestability [16], [17]. In food materials design, these principles must be translated into practical mechanisms. Transparency may include disclosing training data sources, model assumptions, and performance limits. Accountability may require clear lines of responsibility when a model-informed formulation causes quality or safety problems. Contestability means that affected parties, including bakers, workers, and consumers, can challenge decisions that

affect them. Regulatory agencies will need expertise to evaluate computational evidence, especially when models are used to support health, allergen, or labeling claims. The governance challenge is compounded by the fact that food regulations vary across jurisdictions, while supply chains are global.

There is also a need to align computational design with sustainability and energy governance. Machine learning models can be computationally intensive, and their environmental cost must be weighed against their benefits [18], [19]. A model that saves energy in cold storage may justify its training cost, but this should be evaluated rather than assumed. Green AI principles suggest efficiency, reuse, and reporting of computational resources. In frozen bakery systems, the largest environmental impact often comes from refrigeration and transport, so cryoprotectants that reduce temperature sensitivity or extend storage life could have substantial benefits. Governance should incentivize such outcomes while preventing rebound effects, such as increased production enabled by lower waste.

7. Sustainability and Policy Implications

Plant polysaccharide cryoprotectants connect food design to broader sustainability agendas. They can valorize agricultural residues, reduce reliance on synthetic additives, and improve the resilience of cold-chain-dependent food systems. If they reduce freeze-thaw damage, they may lower waste and energy use in distribution. These benefits align with sustainable development goals related to responsible consumption, innovation, and climate action [20]. However, sustainability claims must be assessed across the life cycle. Extraction, purification, drying, and transport can consume energy and water. A polysaccharide derived from a distant side stream may have a higher footprint than a locally sourced alternative. Computational models can support life-cycle assessment by comparing scenarios and identifying hotspots, but they require reliable environmental data.

Policy can shape the direction of innovation. Funding agencies can support open datasets, shared benchmarks, and precompetitive research consortia. Regulatory bodies can develop pathways for evaluating computational evidence and novel food ingredients. Standards organizations can define interoperability requirements for digital twins and cold-chain data. Public procurement can create demand for sustainable frozen bakery products. At the same time, policies must avoid locking in proprietary standards that disadvantage small producers. They should encourage modular architectures, open interfaces, and technology transfer. International coordination is important because frozen bakery supply chains cross borders, and regulatory divergence can slow adoption.

A forward-looking perspective suggests that computational food materials design will mature into an infrastructure discipline. The most valuable assets may not be individual models but the data pipelines, validation protocols, and governance arrangements that make models trustworthy. Plant polysaccharide cryoprotectants will be successful only if they are integrated into this infrastructure. This means designing for manufacturability, cold-chain compatibility, consumer acceptance, and regulatory clarity from the beginning. It also means treating bakers, millers, and consumers as stakeholders rather than passive recipients of optimized formulations. The socio-technical system is not a constraint on design; it is part of the design problem.

8. Conclusion

Computational food materials design of plant polysaccharide-based cryoprotectants offers a powerful pathway for improving frozen bakery systems, but its success depends on system-

level integration. The molecular and colloidal mechanisms of cryoprotection are complex, and plant polysaccharides provide a broad, renewable, and tunable design space. Yet no single model or experiment can capture the full range of interactions that determine dough performance, freeze-thaw robustness, and baking quality. Progress requires coupled simulation, machine learning, data infrastructure, digital twins, and experimental validation. It also requires explicit attention to trade-offs among efficacy, processability, cost, sensory quality, and sustainability.

The deployment of computational cryoprotectants will reshape not only formulations but also relationships among ingredient suppliers, bakeries, equipment manufacturers, regulators, and consumers. Robustness must be designed into models and formulations so that they perform across variable raw materials and cold-chain conditions. Fairness must be considered so that benefits are widely shared and excluded regions or small producers are not disadvantaged. Governance must ensure transparency, accountability, and contestability. Sustainability must be evaluated over the life cycle, not assumed from the biological origin of ingredients. Policy can accelerate responsible innovation by supporting open data, standards, and inclusive participation.

In conclusion, plant polysaccharide cryoprotectants should be understood as components of a socio-technical infrastructure for resilient frozen food systems. Their computational design is most valuable when it is coupled to traceable data, interpretable models, robust process control, and equitable governance. Future research should prioritize integrated platforms that connect molecular structure to industrial deployment, while policy should create incentives for sustainability, openness, and fairness. By treating cryoprotection as a system architecture rather than a single additive function, the field can advance both scientific understanding and practical impact in frozen bakery applications.

References

1. Goff, H. D. (1992). Low-temperature stability and the glassy state in frozen foods. *Food Research International*, 25(4), 317–325. [https://doi.org/10.1016/0963-9969\(92\)90128-R](https://doi.org/10.1016/0963-9969(92)90128-R)
2. Slade, L., & Levine, H. (1991). Beyond water activity: Recent advances based on an alternative approach to the assessment of food quality and safety. *Critical Reviews in Food Science and Nutrition*, 30(2-3), 115–360. <https://doi.org/10.1080/10408399109527543>
3. Selomulyo, V. O., & Zhou, W. (2007). Frozen bread dough: Effects of freezing storage and dough improvers. *Journal of Cereal Science*, 45(1), 1–17. <https://doi.org/10.1016/j.jcs.2006.10.003>
4. Bhattacharya, M., Langstaff, T. M., & Berzonsky, W. A. (2003). Effect of frozen storage and freeze-thaw cycles on the rheological and baking properties of frozen doughs. *Food Research International*, 36(4), 365–372. [https://doi.org/10.1016/S0963-9969\(02\)00228-1](https://doi.org/10.1016/S0963-9969(02)00228-1)
5. Ribotta, P. D., León, A. E., & Añón, M. C. (2003). Effect of freezing and frozen storage on the gelatinization and retrogradation of amylopectin in dough. *Food Research International*, 36(4), 357–363. [https://doi.org/10.1016/S0963-9969\(02\)00227-X](https://doi.org/10.1016/S0963-9969(02)00227-X)
6. Roos, Y. H. (1995). *Phase transitions in foods*. Academic Press.
7. Damodaran, S. (1996). Amino acids, peptides, and proteins. In O. R. Fennema (Ed.), *Food chemistry* (3rd ed., pp. 321–429). Marcel Dekker.

8. Kizilay, E., Kayitmazer, A. B., & Dubin, P. L. (2011). Complexation and coacervation of polyelectrolytes with oppositely charged colloids. *Advances in Colloid and Interface Science*, 167(1-2), 24–37. <https://doi.org/10.1016/j.cis.2011.06.001>
9. Zhang, T., Fang, J. Q., Wang, P. P., & Chen, C. (2026). Structural basis of the cryoprotective sol-gel transition in a *Phyllanthus emblica* L. polysaccharide fraction for frozen dough applications. *Food Hydrocolloids*, 112842.
10. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., & Walsh, A. (2018). Machine learning for molecular and materials science. *Nature*, 559(7715), 547–555. <https://doi.org/10.1038/s41586-018-0337-2>
11. Sanchez-Lengeling, B., & Aspuru-Guzik, A. (2018). Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400), 360–365. <https://doi.org/10.1126/science.aat2663>
12. Ramprasad, R., Batra, R., Pilia, G., Mannodi-Kanakkithodi, A., & Kim, C. (2017). Machine learning in materials informatics: Recent applications and prospects. *npj Computational Materials*, 3, 54. <https://doi.org/10.1038/s41524-017-0056-5>
13. Keith, J. A., Vassilev-Galindo, V., Cheng, B., Chmiela, S., Gastegger, M., Müller, K.-R., & Tkatchenko, A. (2021). Combining machine learning and computational chemistry for predictive insights into chemical systems. *Chemical Reviews*, 121(16), 9816–9872. <https://doi.org/10.1021/acs.chemrev.1c00107>
14. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
15. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
16. Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>
17. Floridi, L., & Cows, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
18. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. <https://doi.org/10.18653/v1/P19-1355>
19. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
20. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Nerini, F. F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, 233. <https://doi.org/10.1038/s41467-019-14108-y>