

Deep Learning-Based Prediction of Dough Rheological Properties Under Repeated Freeze–Thaw Conditions

Gerti Fliegner

Department of Computer Science, University of Houston, Houston, TX, USA.
gertif@uh.edu

Danique Karels

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
hellodanique@oregonstate.edu

Wazir Johal

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
contactwazir@buffalo.edu

Leia Hogberg

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
leia1971@ucf.edu

Abstract

Repeated freeze–thaw conditions impose cumulative structural and rheological changes on dough, altering its viscoelastic behavior, water distribution, gluten network integrity, and fermentation performance. Predicting these changes is difficult because dough is a heterogeneous, history-dependent material whose behavior emerges from coupled biochemical, colloidal, and mechanical processes. This paper examines deep learning-based prediction of dough rheological properties under repeated freeze–thaw cycles from a system-level perspective. It synthesizes literature on frozen dough science, rheological measurement, machine learning architectures, deployment infrastructure, governance, and sustainability. The discussion emphasizes that predictive performance depends not only on model selection but also on sensing quality, experimental design, metadata discipline, uncertainty quantification, and operational integration. Deep learning can combine rheometric time series, imaging, spectroscopy, thermal data, and formulation records, yet such multimodal systems introduce trade-offs in cost, latency, interpretability, and maintainability. The paper argues for hybrid architectures that embed physical priors while retaining representation learning capacity, supported by calibrated uncertainty estimates and transparent documentation. It further considers fairness across flour varieties, regional supply chains, and small-scale bakeries, as well as policy implications for food safety, labeling, data sharing, and energy use. Case illustrations from industrial freezing, artisanal production, cloud kitchens, and humanitarian food systems demonstrate that deployment contexts shape both model requirements and social consequences. The conclusion proposes a research agenda centered on robust benchmarks, causal inference, edge-to-cloud orchestration, and equitable governance for intelligent frozen dough systems.

Keywords

dough rheology; freeze–thaw cycles; deep learning; food systems; uncertainty quantification; infrastructure governance.

1. Introduction

Dough is a complex, time-dependent material whose rheological behavior reflects the interplay of starch granules, gluten proteins, water, yeast, enzymes, lipids, and dissolved solutes. Under repeated freeze–thaw conditions, these components undergo structural reorganization that is difficult to reverse, including ice crystal growth, water migration, gluten depolymerization, and changes in gas retention capacity. Because rheological properties such as extensibility, stickiness, elastic recovery, and resistance to deformation determine processability and final bread quality, predicting them before production is a central challenge for frozen dough manufacturers, retail bakeries, and food service operators. Conventional rheological testing provides accurate measurements but is often slow, sample-intensive, and ill-suited to continuous industrial decision-making. Deep learning has therefore become attractive as a complementary or anticipatory layer that can estimate rheological states from sensor streams, images, and process histories [1], [2].

The promise of deep learning in this domain lies in its ability to model nonlinear and high-dimensional relationships that resist simple parametric description. Dough subjected to repeated freeze–thaw cycles does not degrade linearly, and its rheology can shift abruptly after particular thresholds of ice recrystallization or water redistribution [13], [14]. A prediction system must therefore represent cumulative history rather than a single snapshot. Recurrent neural networks, temporal convolutional models, and attention-based architectures offer mechanisms for learning from sequences of temperature, humidity, mechanical energy input, and time [3], [4]. When combined with imaging or spectroscopy, these models can associate visible or spectral markers with mechanical outcomes, potentially enabling earlier interventions in freezing, thawing, and proofing schedules [17], [20].

From a systems perspective, however, prediction is not merely a modeling problem. It is an infrastructure problem that spans sensors, data pipelines, calibration protocols, model lifecycle management, operator interfaces, and regulatory accountability. A model that performs well on controlled laboratory data may fail in a bakery where freezer doors open frequently, flour batches vary, and ambient conditions fluctuate. The hidden technical debt of machine learning systems, including fragile data dependencies, undeclared feedback loops, and configuration drift, becomes especially consequential when predictions influence food quality and safety [8]. Thus, the design of deep learning systems for dough rheology must confront trade-offs among accuracy, robustness, interpretability, cost, and governance.

This paper develops a system-oriented analysis of deep learning for predicting dough rheological properties under repeated freeze–thaw conditions. It does not propose a single algorithm or report a narrow experimental result. Instead, it examines the layered architecture of such systems, the rheological mechanisms that make prediction difficult, the data and sensing infrastructure required, and the social and policy implications of deploying predictive intelligence in food supply chains. It draws on frozen dough science [13], [14], [15], [16], deep learning foundations [1], [2], food technology reviews [17], [19], [20], and emerging evidence on cryoprotective sol–gel transitions in frozen dough applications [18]. The argument is that successful deployment depends on coupling predictive models with robust measurement, uncertainty-aware decision rules, and transparent governance.

2. Rheological Behavior and Freeze–Thaw Mechanisms

The rheology of dough under repeated freeze–thaw conditions is governed by coupled physical and biochemical transformations. During freezing, water crystallizes into ice, concentrating solutes in unfrozen regions and mechanically disrupting the gluten network. During thawing, melting ice releases water that may not fully reabsorb into the protein matrix, leading to local heterogeneity and altered viscoelastic response [13]. Repeated cycles amplify these effects because ice crystals can recrystallize, grow, and puncture cellular or network structures. The result is a dough that may become stiffer in some deformation regimes, more extensible in others, or more prone to stickiness and tearing depending on formulation and processing history [14], [15].

Rheological measurements capture different manifestations of these changes. Oscillatory tests describe the balance between elastic storage and viscous loss, providing insight into network connectivity and relaxation behavior. Extensional tests reveal resistance to stretching and rupture, which is critical for sheeting, dividing, and molding. Empirical instruments such as farinographs and mixolab systems assess mixing tolerance and pasting behavior, while texture analysis can quantify surface stickiness. Each measurement is sensitive to sample preparation, temperature history, resting time, and deformation rate. Consequently, the same dough can appear stable under one protocol and degraded under another [16]. This protocol dependence complicates the construction of ground truth for machine learning because labels are not absolute material properties but measurement outcomes shaped by instruments and procedures.

Repeated freeze–thaw cycles also interact with fermentation. Yeast cells experience osmotic stress, ice damage, and metabolic perturbation. Their activity after thawing influences gas production, which in turn affects dough expansion and rheological development. If a model predicts rheology only from mechanical and thermal variables while ignoring biological state, it may fail in cases where yeast viability changes independently of gluten structure. Conversely, including fermentation markers such as gas pressure, pH, or volatile compounds can improve prediction but adds sensing complexity [17]. The resulting system must decide which variables are measured online, which are inferred, and which are assumed stable across batches.

Recent work on cryoprotective sol–gel transitions illustrates that certain polysaccharide fractions can modulate ice formation and water mobility in frozen dough systems [18]. Such findings suggest that rheological prediction cannot be purely data-driven in a black-box sense; it must account for formulation chemistry and phase behavior. A deep learning model may learn associations between ingredient descriptors and rheological outcomes, but without physical priors it may generalize poorly to new hydrocolloids, flour blends, or processing regimes. Hybrid approaches that incorporate mechanistic constraints, mass transport concepts, or thermal history summaries can improve extrapolation while retaining the flexibility of neural representation learning [5]. The central design tension is between model capacity and the need for stable behavior under distribution shift.

3. Data Acquisition, Experimental Infrastructure, and Governance

A deep learning system for dough rheology depends on an experimental infrastructure that produces consistent, well-documented data. Rheometers, dynamic mechanical analyzers, imaging systems, near-infrared spectrometers, and thermal sensors each contribute different views of the material. The temporal resolution, spatial resolution, and calibration stability of these instruments determine the upper bound of predictive accuracy. If temperature is measured at the freezer wall rather than at the dough core, the model may learn spurious

correlations. If image lighting changes between batches, a convolutional network may classify shadows instead of surface structure [17], [20]. Data quality is therefore not a preprocessing detail but a foundational system property.

Metadata governance is equally important. Each sample should be accompanied by information about flour provenance, protein content, moisture, additives, mixing energy, resting time, freezing rate, thawing method, cycle count, and ambient conditions. Without such metadata, models cannot distinguish between a formulation effect and a process effect. Datasheets for datasets offer a useful framework for documenting provenance, composition, collection procedures, and recommended uses [10]. In food systems, metadata also intersects with traceability regulations and commercial confidentiality. Manufacturers may be reluctant to share formulation details, while regulators may require evidence of safety and consistency. A predictive system must therefore operate within a data governance regime that balances openness, privacy, and proprietary interests.

The experimental design must also address batch effects and class imbalance. Repeated freeze–thaw degradation is often gradual, but critical failures may occur only after many cycles or under unusual temperature abuse. If the training set contains few examples of severe degradation, the model may underestimate risk in exactly the cases that matter most. Active learning and targeted sampling can help, but they require a feedback loop between prediction, measurement, and model updating. This loop is an organizational process, not merely a technical one. It demands protocols for when to trust predictions, when to recalibrate instruments, and when to override the model with human judgment [8]. Governance thus becomes embedded in daily operations through checklists, audit trails, and decision rights.

4. Deep Learning Architectures for Rheological Prediction

Deep learning architectures for dough rheology can be grouped by the data modalities they address. For imaging and microscopy, convolutional neural networks can identify ice crystal patterns, surface cracks, pore structure, and gluten network discontinuities. For temporal process data, recurrent networks such as long short-term memory models can capture dependencies across freeze–thaw cycles [3]. More recent attention-based models can weigh different time steps and sensor channels, potentially identifying critical windows such as the initial freezing phase or the final thaw [4]. For spectral data, one-dimensional convolutional networks and fully connected encoders can learn features associated with water binding, protein secondary structure, and lipid oxidation.

Multimodal fusion is particularly relevant because no single sensor fully describes dough state. Early fusion concatenates features before prediction, while late fusion combines separate predictions. Intermediate fusion can learn cross-modal interactions, such as how a spectral water marker modifies the interpretation of a mechanical relaxation curve. However, fusion increases data requirements and model complexity. Missing modalities are common in industrial settings, so architectures must handle incomplete inputs without catastrophic failure. This requirement favors modular designs with fallback pathways and uncertainty estimates that increase when a modality is absent [6]. A system that predicts confidently from incomplete data may be more dangerous than one that abstains.

Physics-informed neural networks offer a way to embed domain constraints into learning [5]. In dough rheology, such constraints might include monotonicity of degradation under repeated temperature abuse, conservation of water mass across phases, or bounded relaxation behavior. These constraints can be implemented through architecture choices, data

augmentation, or penalty terms that discourage physically impossible predictions. The advantage is improved generalization and interpretability; the disadvantage is that incorrect physical assumptions can bias predictions. A hybrid system might use physics-informed components for well-understood processes and data-driven components for unresolved mechanisms such as hydrocolloid interactions or yeast stress responses.

Transfer learning and foundation models are also relevant. A model pretrained on general food images or spectral data may require less labeled dough data than a model trained from scratch. Yet transfer from unrelated domains can introduce negative transfer if the learned features are misaligned with rheological mechanisms. Domain adaptation techniques can adjust representations across flour types, equipment, or geographic regions, but they require some target-domain data. In practice, deployment may involve a global model pretrained on diverse data and local fine-tuning at a bakery or factory. This architecture shifts governance questions: who owns the local adaptation, how are updates validated, and how are disparities across sites detected [9], [10].

5. System Architecture and Deployment Trade-Offs

The deployment architecture for dough rheology prediction must balance latency, bandwidth, privacy, and resilience. Cloud-based inference enables large models, centralized monitoring, and frequent updates, but it depends on reliable connectivity and introduces data transfer costs. Edge inference on local gateways or programmable logic controllers reduces latency and keeps sensitive production data on site, but it constrains model size and complicates updates. A hybrid architecture may run lightweight anomaly detection at the edge and send uncertain or high-stakes cases to the cloud for deeper analysis. This tiered design reflects a broader principle in cyber-physical systems: computation should be placed where it best matches decision timing and data sensitivity [8], [19].

Integration with existing manufacturing execution systems is another critical layer. If predictions are displayed in a separate dashboard that operators ignore, the system has no operational effect. If predictions automatically adjust freezer settings, they must be validated for safety and failure modes. Human-in-the-loop design is not a concession to immature technology but a governance mechanism that preserves accountability. Operators need explanations that are meaningful in their work context, such as expected changes in dough handling rather than abstract feature importance. Model cards can document intended use, limitations, and performance across subgroups, supporting both operators and auditors [9]. Datasheets for datasets complement this by describing training data provenance and potential biases [10].

Maintenance and lifecycle management are often underestimated. Sensors drift, ingredients change, and freezer performance degrades. A model trained on last year's flour may become miscalibrated after a harvest transition. Monitoring systems should track input distributions, prediction distributions, and downstream quality outcomes. When drift is detected, the organization must decide whether to recalibrate sensors, collect new labels, fine-tune the model, or revert to a previous version. This decision process requires versioning, rollback capability, and clear ownership. Without such infrastructure, predictive accuracy will decay silently, and operators may lose trust after a few visible failures [8]. The technical debt of machine learning systems is thus directly connected to food quality risk.

Cost is a cross-cutting constraint. High-end rheometers, imaging systems, and spectral sensors may be affordable for large manufacturers but not for small bakeries. A system designed only

for well-resourced facilities can widen productivity gaps and reduce fairness across the food sector. Lower-cost sensing alternatives, such as smartphone cameras, inexpensive temperature loggers, and simplified mechanical probes, may provide sufficient signal for many decisions. However, they introduce calibration and robustness challenges. A tiered product ecosystem could offer basic predictions for small operators and advanced analytics for industrial users, but the underlying governance should ensure that safety-critical claims are not made beyond validated conditions [19].

6. Robustness, Uncertainty, and Fairness

Robustness in dough rheology prediction means maintaining performance under distribution shift. Shift can arise from wheat variety, harvest year, milling extraction rate, water hardness, additive substitution, equipment differences, and ambient climate. A model trained in a temperate climate may fail in a tropical bakery where thawing occurs at higher humidity. Robustness strategies include domain randomization during training, adversarial validation, and ensemble methods. Yet robustness is not absolute. The system should know when it is outside its validated domain and communicate that uncertainty to users. Bayesian deep learning provides a principled framework for separating aleatoric uncertainty, which arises from measurement noise, and epistemic uncertainty, which reflects limited knowledge [6]. In food operations, this distinction matters because the response to noise may be repeated measurement, while the response to knowledge gaps may be expert consultation or conservative process settings.

Calibration is central to trustworthy uncertainty. A model that reports high confidence but is frequently wrong can cause unsafe decisions, while an overly uncertain model may be ignored. Modern neural networks are often poorly calibrated, requiring temperature scaling, ensembling, or Bayesian approximations [7]. Calibration should be evaluated not only overall but also across flour types, cycle counts, and equipment classes. A model may be well calibrated on average yet systematically overconfident for high-protein flours or for late freeze–thaw cycles. Such disparities can lead to unequal risk. If small bakeries use cheaper sensors or different freezing practices, they may experience higher error rates even when using the same model. Fairness in this context is therefore socio-technical: it concerns the distribution of predictive reliability and the capacity to act on predictions.

Governance mechanisms should make these issues visible. Model cards can report performance disaggregated by relevant groups, intended use, and out-of-scope conditions [9]. Datasheets can document how data were collected and whether certain regions, flour varieties, or production scales are underrepresented [10]. Audit trails can record when predictions were overridden and why. These instruments are not merely bureaucratic. They create organizational memory that supports learning and accountability. In food safety, where decisions may affect allergens, microbial growth, or labeling claims, transparency is essential. However, transparency must be balanced against proprietary information and the risk that overly complex documentation will not be used. Effective governance therefore requires concise, role-specific reporting.

Another fairness dimension concerns labor. Predictive systems may deskill some tasks by transferring judgment from bakers to models, while creating new tasks in data curation and exception handling. If workers are evaluated against model recommendations without understanding their limitations, trust and morale can suffer. Participatory design can help align model outputs with shop-floor knowledge. Bakers often detect dough changes by touch, sound, and smell before sensors register them. A deep learning system should augment rather

than replace these embodied skills. The most robust architecture may be one that treats human observation as an additional sensor channel, recorded and contextualized rather than dismissed as subjective [12].

7. Sustainability, Governance, and Policy Implications

Frozen dough systems are embedded in energy-intensive cold chains. Freezing, storage, transport, and thawing consume electricity and refrigerants, while inefficient prediction can lead to unnecessary waste or overproduction. Deep learning can contribute to sustainability by improving process control, reducing rejected batches, and enabling dynamic freeze–thaw scheduling. If a model predicts that a dough batch will remain within acceptable rheological limits for an additional cycle, a manufacturer may avoid premature disposal. If it predicts rapid degradation, the batch can be redirected to products with lower rheological demands. These decisions can reduce food waste and energy use. However, the environmental cost of training and deploying large models must also be considered. Energy and policy considerations for deep learning show that model size and training frequency have material environmental implications [11]. Sustainable design favors efficient architectures, edge inference, and retraining schedules based on need rather than fashion.

Policy frameworks for food safety and labeling are not yet fully adapted to predictive analytics. If a model influences process decisions, regulators may ask whether it is validated, whether it is biased, and whether it can be audited. Standards for software in food production exist, but they often focus on deterministic control systems rather than adaptive machine learning. A predictive model that updates continuously may fall between traditional validation and modern MLOps practices. Policymakers could encourage risk-based validation, where the level of scrutiny depends on the consequences of error. For low-risk decisions such as display recommendations, lighter oversight may suffice. For safety-critical decisions such as allergen cross-contact or fermentation control, stronger evidence and human oversight are needed.

Data sharing is another policy tension. Collaborative learning across bakeries, mills, and equipment vendors could improve model robustness, especially for rare failure modes. Yet formulation and process data are commercially sensitive. Federated learning and privacy-preserving analytics offer technical possibilities, but they require governance agreements about ownership, access, and liability. Public research infrastructures could provide neutral benchmarks and anonymized datasets, similar to open datasets in computer vision and agriculture [19]. Such resources would allow independent evaluation of claims and reduce the risk that proprietary models encode hidden biases. They would also support reproducibility, which is currently limited in food machine learning because datasets and preprocessing pipelines are often unpublished [17], [20].

The broader policy question is whether intelligent frozen dough systems will centralize or distribute value. Large manufacturers may benefit most from data-driven optimization, while small bakeries may struggle to adopt and maintain complex systems. Public investment in low-cost sensors, open standards, and extension services could democratize access. Procurement policies could require interoperability and data portability, preventing vendor lock-in. Training programs could help workers interpret model outputs and participate in governance. Without such measures, predictive systems may improve efficiency in some segments while deepening inequalities in others. The policy goal should be resilient, equitable food systems rather than narrow automation.

8. Case Illustrations and Cross-Domain Comparisons

Consider an industrial frozen dough plant producing laminated pastry and bread rolls. Freeze–thaw cycles occur not only during initial freezing but also during transport disruptions, storage defrost cycles, and thawing before shaping. The plant may have inline temperature sensors, imaging for surface defects, and periodic rheometer samples. A deep learning system could fuse these streams to predict extensibility before sheeting. The main constraints are high throughput and cost of false alarms. If the model predicts failure too often, operators will ignore it. If it misses failure, production lines may jam or final products may be rejected. The architecture must therefore prioritize calibrated thresholds and fast inference, likely at the edge, with cloud-based retraining. Governance must define who can adjust thresholds and how overrides are documented [8], [9].

An artisanal bakery presents a different context. Freezing may be used to batch production, but equipment is less controlled and batches are small. A smartphone-based imaging system and a simple temperature logger might be the only affordable sensors. A lightweight model could predict whether thawed dough is suitable for hand shaping or should be rested longer. Here, interpretability and trust matter more than raw accuracy. The baker needs actionable guidance, not a probabilistic score alone. The model should account for local flour and climate, which may not be represented in global training data. A transfer learning approach with local fine-tuning could help, but it requires a user-friendly way to label outcomes. This case illustrates that deployment context determines the optimal trade-off between model complexity and maintainability [19].

Cloud kitchens and food delivery operations face volatile demand and frequent freezing of portioned dough. Prediction can support dynamic thawing schedules that match expected orders, reducing waste while maintaining quality. The system may integrate with order forecasting, inventory, and equipment control. This integration raises fairness questions about labor scheduling and algorithmic management. If predictions drive worker tasks without transparency, they may intensify pace and reduce autonomy. A more equitable design would give workers visibility into predictions and the ability to flag discrepancies. It would also avoid using rheological predictions as covert performance metrics. The technical system is thus entangled with workplace governance [12].

Humanitarian food systems provide a contrasting case where sustainability and equity are paramount. In refugee camps or disaster responses, frozen dough may be used for centralized baking, but energy and refrigeration may be unreliable. A predictive model that requires constant connectivity is inappropriate. Instead, simple thermal indicators and rule-based decision support may be more robust. Deep learning could still contribute through offline models on low-power devices, but validation must account for extreme variability in ingredients and conditions. Cross-domain comparisons with precision agriculture are instructive: both domains involve biological variability, environmental uncertainty, and decisions that affect livelihoods [19]. In agriculture, deep learning has improved detection and yield prediction but has also raised concerns about data ownership and access. Frozen dough systems should learn from those debates rather than repeat them.

9. Limitations and Future Research Directions

Current research on deep learning for dough rheology remains fragmented. Many studies use small datasets from a single laboratory, lack external validation, and report aggregate accuracy without uncertainty or subgroup analysis. This limits claims about generalizability. Future work should establish open benchmarks with standardized metadata, multiple flour types, controlled freeze–thaw protocols, and held-out facilities. Benchmarks should include

not only prediction error but also calibration, robustness to missing modalities, inference latency, and energy consumption. They should be accompanied by datasheets and model cards so that comparisons are meaningful [10]. Without such infrastructure, the field risks producing models that perform well in publication but poorly in practice.

Causal inference is another underdeveloped area. Predictive models learn correlations between sensor signals and rheological outcomes, but interventions require causal understanding. If a model suggests lowering freezer temperature to preserve extensibility, it should be based on evidence that the change causes the desired effect rather than merely correlating with it. Randomized experiments or natural experiments across production lines could support causal claims. Hybrid models that combine physics with causal graphs may be more useful for decision support than purely predictive models [5]. They can also help identify confounding variables, such as flour quality, that affect both sensor readings and rheological outcomes.

Human factors and governance deserve more attention. Research on model cards, datasheets, and technical debt provides useful starting points [8], [9], [10], but food systems have specific constraints related to safety, labor, and regulation. Future studies should examine how operators interpret uncertainty, how auditors evaluate adaptive models, and how small enterprises can participate in data collaboratives. Participatory design methods could reveal mismatches between model outputs and shop-floor decision-making. Longitudinal studies could track trust and performance over multiple harvests and equipment changes. Finally, sustainability assessment should include not only direct energy use but also food waste, packaging, and cold-chain infrastructure. The required transition toward resilient food systems will depend on whether predictive intelligence is designed as a public good or a proprietary advantage.

10. Conclusion

Deep learning offers a powerful but partial solution to predicting dough rheological properties under repeated freeze–thaw conditions. The material is history-dependent, heterogeneous, and sensitive to formulation and process variables. Models can learn complex mappings from sensor data to rheological outcomes, especially when they combine temporal, spectral, and imaging modalities. Yet predictive accuracy alone is insufficient. Deployment requires robust sensing, disciplined metadata, uncertainty quantification, edge-to-cloud architecture, and clear governance. Systems that ignore these layers may fail silently, erode trust, or exacerbate inequities across the food sector.

The most promising path is hybrid and human-centered. Physics-informed components can encode known mechanisms, while data-driven components capture unresolved interactions. Calibrated uncertainty can support abstention and expert consultation. Model cards and datasheets can make performance and provenance visible. Policy frameworks can encourage open benchmarks, interoperability, and equitable access to low-cost sensing. By treating dough rheology prediction as a socio-technical infrastructure rather than a standalone algorithm, researchers and practitioners can improve quality, reduce waste, and support more sustainable frozen dough systems. The agenda is ambitious, but the convergence of food science, machine learning, and governance makes it both necessary and feasible.

References

1. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
<https://doi.org/10.1038/nature14539>

2. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
3. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
5. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
6. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 30, 5574–5584.
7. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. *Proceedings of the 34th International Conference on Machine Learning*, 70, 1321–1330.
8. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
9. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. <https://doi.org/10.1145/3287560.3287596>
10. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
11. 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>
12. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
13. 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>
14. Baier-Schenk, A., Handschin, S., von Schönau, M., Bittermann, A. G., Báez, L., & Conde-Petit, B. (2005). In situ observation of the freezing process in wheat dough by confocal laser scanning microscopy: Formation of ice crystals and changes in the gluten network. *Journal of Cereal Science*, 42(2), 181–187. <https://doi.org/10.1016/j.jcs.2005.04.001>

15. Räsänen, J., Blanshard, J. M. V., Mitchell, J. R., Derbyshire, W., & Autio, K. (1998). Properties of frozen wheat doughs at subzero temperatures. *Journal of Cereal Science*, 28(1), 1–14. <https://doi.org/10.1006/jcrs.1997.0159>
16. Kenny, S., Wehrle, K., Auty, M. A. E., & Arendt, E. K. (2000). Influence of sodium caseinate and whey protein on baking properties and rheology of frozen dough. *Cereal Chemistry*, 77(3), 365–370. <https://doi.org/10.1094/CCHEM.2000.77.3.365>
17. Zhou, L., Zhang, C., Liu, F., Qiu, Z., & He, Y. (2019). Application of deep learning in food: A review. *Comprehensive Reviews in Food Science and Food Safety*, 18(6), 1793–1811. <https://doi.org/10.1111/1541-4337.12492>
18. 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.
19. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90. <https://doi.org/10.1016/j.compag.2018.02.016>
20. Zhu, L., Spachos, P., Pensini, E., & Plataniotis, K. N. (2021). Deep learning and machine vision for food processing: A survey. *Current Research in Food Science*, 4, 233–249. <https://doi.org/10.1016/j.crfs.2021.03.009>