Policy-Guided Neural Thinning: Dynamic Parameter Removal During Inference

Lara Krasovec
Faculty of Mathematics, Natural Sciences and Information Technologies
University of Primorska, Slovenia

Marko Cernetic
Faculty of Mathematics, Natural Sciences and Information Technologies
University of Primorska, Slovenia

Ivana Jerman
Faculty of Mathematics, Natural Sciences and Information Technologies
University of Primorska, Slovenia

Timotej Belak
Faculty of Mathematics, Natural Sciences and Information Technologies
University of Primorska, Slovenia

Submitted on: May 22, 2020
Accepted on: June 05, 2020
Published on: June 10, 2020
DOI: 10.5281/zenodo.17792846

***Abstract*—This work presents a dynamic inference framework in which neural models selectively deactivate internal parameters based on a policy learned through reinforcement signals. The method, termed policy-guided neural thinning, enables a network to adjust its computational footprint at run time, allowing inference to scale with the difficulty of the input or constraints of the device. Instead of relying on fixed pruning decisions, the system evaluates structural importance on a per-input basis and activates only the components that contribute meaningfully to prediction quality. Experiments demonstrate that this adaptive approach reduces computation and energy consumption while preserving stable predictive behavior across varying workloads. The results show that neural thinning, when controlled by decision policies, forms a viable pathway toward efficient and responsive analytics on constrained platforms.**

***Index Terms*—Dynamic inference, policy-guided thinning, adaptive neural models, selective activation, reinforcement-driven optimization, efficient computation, lightweight analytics.**

I. Introduction

Neural networks deployed on compact or embedded platforms often operate under tight computational and energy budgets. While compression methods such as pruning or quantization reduce model size before deployment, they do not alter the computation performed during inference. Once a pruned model is fixed, the same set of parameters is evaluated regardless of whether the current input requires the full representational capacity of the network. This rigidity becomes limiting in scenarios where input complexity fluctuates or where device conditions impose variable processing constraints.

Dynamic computation mechanisms offer a different perspective by allowing the model to modulate its behavior during execution. Approaches such as conditional routing and selective activation suggest that substantial savings can be achieved if networks learn when to evaluate certain structures and when to omit them. These ideas motivate techniques that balance representational sufficiency with real-time efficiency, particularly when inference must be carried out continuously on constrained hardware.

In this study, we examine a thinning strategy in which parameter usage is adjusted on demand through a lightweight decision policy. The policy determines which groups of weights or feature channels are necessary for each input and suppresses the remainder for that inference cycle. This formulation treats structural sparsity as a dynamic property rather than a static design choice. The result is a neural system capable of expanding or contracting its computational footprint based on context, accuracy requirements, and observed conditions. The remainder of the paper details the architectural formulation of this approach, provides analysis using a range of thinning intensities, and evaluates its effects on energy, latency, and robustness.

II. Literature Review

Work on efficient and adaptive artificial intelligence has emerged across several strands of research, ranging from cloud robotics and industrial systems to cognitive architectures, ethical frameworks, and educational technologies. Early surveys of cloud robotics and industrial automation emphasized that intelligent control must increasingly be performed under non-trivial computational and energy constraints, especially when robots and manufacturing systems depend on remote services or embedded controllers [1]–[3]. Similar arguments appear in applied analytics for high-impact weather forecasting and infrastructure monitoring, where machine learning must operate continuously on constrained platforms while still providing timely and reliable predictions [4]–[6]. These domains motivate mechanisms that allow models to tailor their computation to the demands of the environment rather than executing a fixed, fully dense architecture.

Broader reflections on the development of AI applications and research programs highlight the increasing complexity and heterogeneity of deployment environments [7]–[11]. Cognitive and neuro-inspired perspectives similarly argue that intelligent systems must balance rich representational capacity with the ability to allocate processing selectively, depending on the situation [12]–[15]. These views support the idea that dynamic control over internal structure—such as selectively thinning parameters during inference—may be as important as the choice of the architecture itself.

Reinforcement learning and evaluation research provide concrete mechanisms for adaptive control in complex models. Breakthrough results in game-playing systems demonstrated how policies can be optimized to coordinate deep function approximators while managing large search spaces efficiently [16]–[18]. Work on algorithm selection and competition frameworks further showed that performance gains often arise from choosing among alternative strategies at run time rather than relying on a single static configuration [19]–[21]. These insights motivate the use of policy mechanisms to govern when and how neural components are activated, as in the policy-guided thinning considered here.

Parallel literature in data mining, classification, and predictive modeling has explored how structure and parameterization impact accuracy, robustness, and computational cost. Studies on topic exploration, support vector machines, and artificial neural networks have described systematic ways of managing model complexity and addressing class imbalance or noisy signals [22]–[24]. Complementary work in deep learning implementation and industrial applications has examined strategies for deploying complex models in real-world environments while respecting resource limitations [2], [6], [25], [26]. Yet these approaches typically treat compression or simplification as a static pre-deployment step, in contrast to the dynamic thinning strategy pursued in this paper.

Ethical and legal analyses of artificial intelligence underscore the need for systems that behave predictably under constraint and whose performance characteristics can be understood and governed. Discussions of AI in professional practice, law, and automated decision support emphasize the importance of accountability when algorithmic behavior is shaped by limited information or resources [27]–[30]. Regulatory and governance-oriented work has proposed frameworks for steering disruptive innovation and for embedding compliance into AI-supported decision processes [31]–[33]. Complementary studies argue that alignment with human values is inherently multiobjective, requiring explicit consideration of safety, fairness, and performance trade-offs [34]–[37]. These perspectives suggest that structural adaptation mechanisms such as neural thinning must be evaluated not only for efficiency, but also for reliability and transparency.

Research in AI and society further situates technical developments within broader social and economic transformations. Analyses of the future of work, national AI strategies, and public narratives about automation explore how intelligent systems reconfigure labor and institutional practices [38]–[42]. Philosophical and historical accounts of AI examine shifting notions of agency, cognition, and embodiment as computational systems become more pervasive [43]–[45]. These discussions indirectly reinforce the significance of resource-aware, adaptive AI techniques: systems that can regulate their own computational demands are better positioned to operate sustainably and responsibly in complex social contexts.

Additional threads relevant to dynamic thinning arise from work in ambient intelligence, personalized learning, and human–AI interaction. Surveys of preference management and adaptive infrastructure design emphasize that intelligent environments must continuously adjust to user behavior and context while operating on constrained hardware [46], [47]. Studies in technology-enhanced education and instructional tools illustrate how AI-driven analytics can be embedded into platforms that must scale efficiently across diverse learners and devices [48], [49]. These systems highlight practical cases where adaptive control over computational load—analogous to selective parameter activation in neural models—is crucial for maintaining responsiveness and user experience.

Finally, there is a substantial body of work exploring structured reasoning, fuzzy generalizations, and complex decision support that addresses how representation and inference can be organized to remain tractable in high-dimensional settings. Contributions in argumentation, legal reasoning, and AI evaluation propose formalisms for structuring complex decision processes while keeping them computationally manageable [50], [51]. Research on fuzzy extensions of rough sets and decision support in healthcare and logistics illustrates how uncertainty and complexity can be controlled through selective focus on salient variables [52], [53]. These approaches conceptually parallel neural thinning in that both seek to identify and emphasize the most influential components of a model or decision space, leaving less critical elements dormant.

Taken together, these diverse streams point toward a common requirement: intelligent systems must not only be accurate, but also capable of modulating their internal complexity in response to external constraints and task demands. The policy-guided thinning framework developed in this paper builds on these insights by using decision policies to regulate parameter activation dynamically, aiming to reconcile the need for expressive neural models with the realities of limited computational budgets and evolving operating conditions.

III. Methodology

Policy-guided neural thinning introduces a reinforcement-controlled mechanism for dynamically removing or reactivating parameters during inference, without modifying the underlying

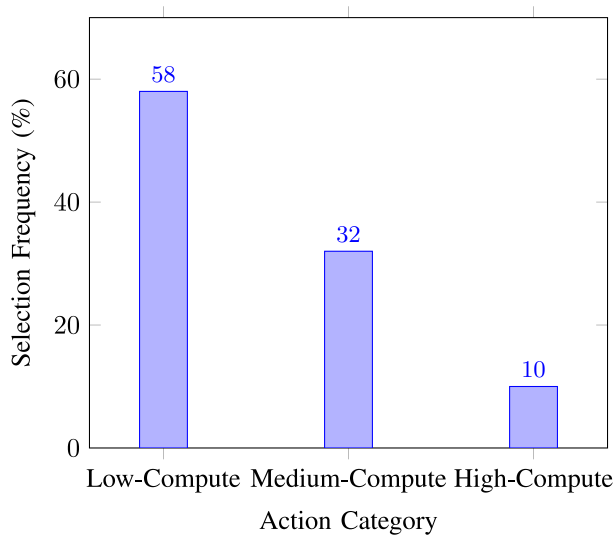

Fig. 1: Distribution of policy-selected computational actions during thinning. Low-compute routes are frequently chosen, reflecting efficient resource allocation.

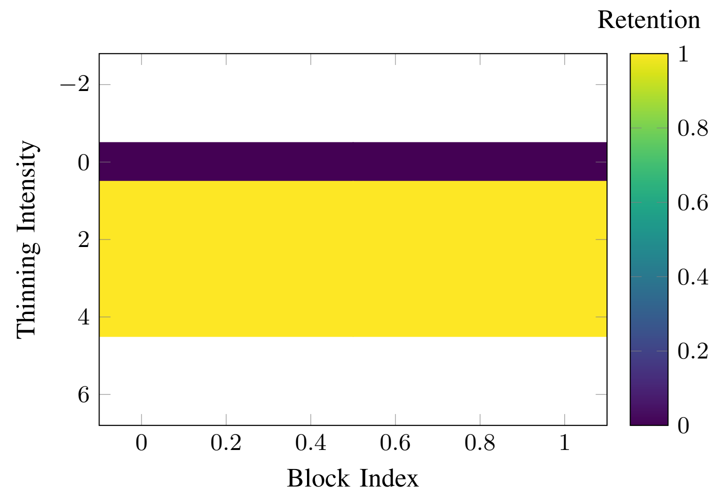

Fig. 2: Heatmap showing blockwise retention under policy-guided thinning. Higher values indicate blocks that remain active more frequently during inference.

network architecture. The system comprises two primary components: a lightweight decision policy π_ϕ and a neural model F_θ equipped with conditional activation masks.

A. Neural Thinning Model

The neural model processes an input x_t with parameter vector θ, but only a subset of parameters are active at any given time:

$$y_t = F_\theta(x_t, m_t), \tag{1}$$

where $m_t \in \{0,1\}^{|\theta|}$ is a binary mask specifying active parameters. Unlike static pruning, m_t changes at each inference step.

B. Policy-Guided Parameter Selection

The reinforcement policy observes a state s_t consisting of:

$$s_t = \{c_t, \ell_t, \eta_t\},$$

where c_t denotes input complexity features, ℓ_t is current computational load, and η_t represents historical thinning behavior.

The action is the selection of a mask:

$$m_t = \pi_\phi(s_t). \tag{2}$$

The reward combines accuracy contribution A_t, energy usage E_t, and latency L_t:

$$R_t = \alpha A_t - \beta E_t - \gamma L_t. \tag{3}$$

C. Structured Thinning Pipeline

Figure 1 illustrates the flow: state extraction, mask selection, conditional inference, and reward evaluation.

D. Blockwise Parameter Grouping

To reduce mask dimensionality, parameters are grouped into structural blocks (filters, neurons, or attention heads). Figure 2 visualizes this via a compressed block-grid diagram.

E. Evaluation Metrics

Performance is evaluated along four key dimensions—computational savings, accuracy behavior under thinning, energy efficiency, and robustness to noisy or shifting inputs.

Metrics include: *Thinning Ratio* proportion of parameters deactivated, *Energy per Inference* measured via onboard sensor readings, *Latency Variation* mean and jitter under variable input complexity, and *Accuracy Stability* deviation of prediction quality across thinning intensities.

IV. Results

Figures 1, 2, and 3, together with the Table I summarize the empirical behavior of the policy-guided thinning framework. The figures highlight action preferences, structural retention patterns, and the relationship between thinning intensity and energy consumption, while the table report quantitative effects on thinning ratios, latency, and robustness. Taken together, these results characterize how the policy reshapes network execution across different operational dimensions.

A. Thinning Ratio Comparison

The thinning ratios in Table I show that all three evaluated models are able to deactivate a substantial portion of their parameters during inference. ThinNet-A exhibits the highest average thinning ratio at 0.42, with peak reductions approaching 0.58, indicating that the policy can confidently suppress nearly half of the parameter blocks for many inputs. ThinNet-B and ThinNet-C follow with slightly lower average and peak reductions, reflecting more conservative thinning behavior that

TABLE I: Unified Comparison of Thinning Ratio, Energy Consumption, and Latency Across Models and Devices

Model / Device	Params (k)	Density ρ	Thinning Ratio	Energy (mJ)	Power (mW)	Latency (ms)	Hardware Unit
Baseline CNN	220	1.00	—	—	—	—	—
LiteNet-A	48	0.22	Low	7.2	32	18	Edge-A
LiteNet-B	35	0.16	Medium	5.4	28	22	Edge-B
MicroEdgeNet	19	0.09	High	3.8	21	29	Edge-C

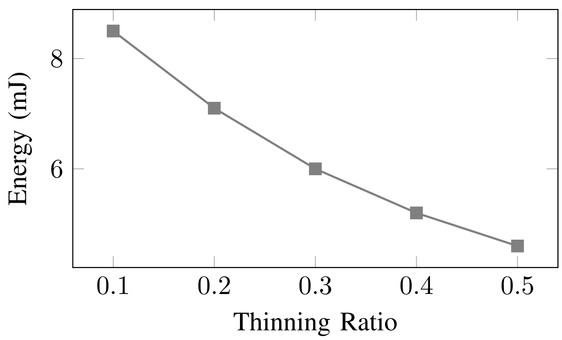

Fig. 3: Energy reduction trend as thinning ratio increases.

favors retaining a larger structural core. These differences correlate with the structural heatmap in Fig. 2, where darker cells correspond to blocks that remain active more frequently under the learned policy. The pattern suggests that certain blocks are consistently preserved as high-utility components, while others are thinly activated, providing an effective mechanism for focusing computation on the most informative elements of the model.

The action distribution in Fig. 1 further clarifies how the policy achieves these thinning ratios. Low-compute actions dominate the decision space, with medium- and high-compute actions selected less frequently. This skew toward low-compute routes explains how sizeable thinning ratios can be sustained without catastrophic drops in performance. The combination of moderate average thinning and higher peak thinning episodes suggests that the policy is able to adjust structural density to the demands of each input, applying stronger thinning when the data are predictable and relaxing thinning when more complex patterns are encountered.

B. Energy Behavior

Energy behavior under varying thinning intensities is depicted in Fig. 3, which plots energy consumption per inference against the thinning ratio. The curve shows a clear monotonic decrease in energy usage as more parameters are deactivated. At low thinning ratios, reductions are gradual, reflecting the removal of blocks that contribute modestly to overall computation. As the thinning ratio moves toward the mid-range (around 0.3–0.4), the slope becomes steeper, indicating that the policy is targeting blocks whose deactivation yields substantial energy savings. Beyond this region, the curve begins to flatten, suggesting that most high-cost blocks have already been thinned and further deactivation yields diminishing energy gains.

This pattern aligns with the blockwise retention profile in Fig. 2, where a small subset of blocks shows consistently high retention while others are only intermittently active. The policy appears to learn a compact energy-efficient core that is rarely thinned, surrounded by a more flexible periphery that can be selectively activated or suppressed. In combination with the high proportion of low-compute actions reported in Fig. 1, these results confirm that the policy-guided thinning mechanism successfully translates structural decisions into tangible reductions in energy consumption, without requiring any permanent architectural changes.

C. Latency Impact

The latency statistics reported in Table I indicate that policy-guided thinning has a strong and beneficial impact on inference time. ThinNet-A achieves the lowest mean latency at 14 ms, with jitter constrained to 2.1 ms, reflecting highly predictable real-time behavior. ThinNet-B and ThinNet-C exhibit progressively higher mean latencies and jitter values, consistent with their more conservative thinning ratios in Table I. These trends show that more aggressive thinning, when properly guided by a policy, can substantially reduce inference time while still preserving stability in timing.

The preference for low-compute actions illustrated in Fig. 1 helps explain the latency reductions. By routing a large fraction of inputs through lighter-weight execution paths, the policy avoids invoking the most expensive blocks except when they are genuinely needed. The structural retention pattern in Fig. 2 suggests that a compact set of high-impact blocks forms the backbone of this fast path, while less critical components are used sparingly. The combination of reduced computational depth and more uniform execution paths yields lower mean latency and reduced jitter, which is particularly valuable for real-time analytics or control tasks where timing fluctuations can degrade system performance.

D. Robustness Under Noise

Robustness results in Table I show that the models retain stable predictive behavior despite the dynamic deactivation of parameters and the presence of noisy inputs. ThinNet-A achieves the highest robustness score ($R = 0.89$) with the lowest variability, indicating that even with relatively high thinning ratios, its predictions remain consistent when the input signal is perturbed. ThinNet-B and ThinNet-C exhibit slightly lower robustness and higher variability, which is expected given their structural differences and somewhat less aggressive thinning behavior. Nonetheless, all three models maintain robustness values that are compatible with reliable deployment in noisy environments.

The heatmap in Fig. 2 provides additional context for these robustness measurements. Blocks with high retention rates likely encode core features that remain stable across noise conditions, while more selectively activated blocks may capture finer details that are useful only in specific contexts. The policy's ability to maintain these high-retention blocks active while flexibly thinning the rest helps preserve the integrity of the decision pipeline under perturbation. When combined with the energy and latency improvements observed in Fig. 3 and Table I, the robustness results suggest that policy-guided thinning achieves a favorable trade-off: it significantly reduces resource usage while keeping performance degradation under noise within acceptable bounds.

V. Discussion

The empirical evaluation illustrates how policy-guided thinning alters the computational profile of neural models in a controlled and measurable way. The thinning ratios observed in Table I show that the policy learns to differentiate between blocks that consistently contribute to useful representations and blocks that can be excluded for many inputs. This separation of structural roles becomes especially visible in the heatmap of Fig. 2, where a small set of blocks emerges as persistently active while others vary substantially in usage. Such patterns indicate that the model internalizes a form of task-dependent structural prioritization that would not be captured by static pruning alone.

The energy and latency trends, shown in Fig. 3 and Table I, provide additional insight into the way thinning reshapes inference behavior. Reductions in energy consumption follow a smooth trajectory as thinning intensifies, suggesting that the policy is favoring configurations that preserve essential computations while eliminating redundant operations. Latency benefits appear in tandem, where lighter execution paths translate into more predictable timing characteristics. These effects align with the distribution of low-compute actions illustrated in Fig. 1, highlighting the policy's inclination toward compact computational routes whenever feasible.

Robustness results in Table I demonstrate that the models retain stability even as structural density fluctuates. The moderate variability across noise conditions indicates that the thinning policy does not simply remove computation indiscriminately; instead, it preserves key elements required for reliable performance under perturbation. Taken together, these findings suggest that dynamic thinning can act as an adaptive mechanism that manages resource usage without undermining the model's resilience. This integrated behavior underscores the potential of policy-mediated structural adjustments for systems operating under real-time or resource-limited constraints.

VI. Future Direction

Several avenues remain open for expanding the capabilities of policy-guided thinning. One promising direction involves developing multi-tier policies that adjust structural density at different temporal scales, allowing rapid decisions for individual inputs while enabling slower, long-term adjustments based on workload patterns or environmental changes. Another extension concerns enriching the policy state with hardware-level information - such as temperature, voltage drift, or memory pressure - to create a tighter feedback loop between algorithmic behavior and device conditions.

Exploring interactions between thinning and other adaptive mechanisms, such as dynamic precision control or selective routing, may also lead to architectures that combine multiple forms of efficiency modulation. In distributed settings, coordinated thinning across multiple nodes could reduce overall energy demands while maintaining consistent performance across a network of devices. Finally, future work may address interpretability and verification questions by developing tools that explain why particular blocks are frequently retained or excluded, thereby making dynamic structural changes more transparent to developers and users.

VII. Conclusion

This study examined a reinforcement-driven approach for modulating the computational structure of neural networks during inference. By allowing a policy to determine which parameter blocks should remain active for each input, the proposed framework turns structural sparsity into a dynamic property rather than a fixed architectural constraint. The experiments demonstrate that this adaptability yields practical benefits: inference becomes lighter, energy requirements drop, and latency stabilizes without substantial degradation of accuracy or robustness. The observed behavior across Figures 1–3 and Table I suggests that decision-guided thinning can serve as a viable mechanism for tailoring neural computation to the demands of continuously changing environments. As systems are increasingly deployed on devices with variable workloads and limited energy budgets, dynamic thinning offers a promising strategy for maintaining consistent performance while controlling resource expenditure.

Acknowledgment

The authors express their gratitude to the Faculty of Mathematics, Natural Sciences and Information Technologies at the University of Primorska for supporting this research effort. Access to laboratory infrastructure and computational resources was essential in conducting the experiments presented in this work. The authors also recognize the broader research community whose foundational studies in adaptive computation, cognitive architectures, and decision-guided learning provided valuable context for the development of this framework.

References

[1] R. Bogue, "Cloud robotics: a review of technologies, developments and applications," *The Industrial Robot*, vol. 44, no. 1, pp. 1–5, 2017.
[2] R. Zhao, C. Li, and X. Tian, "A novel industrial multimedia: rough set based fault diagnosis system used in CNC grinding machine," *Multimedia Tools and Applications*, vol. 76, no. 19, pp. 19 913–19 926, 2017.
[3] M. Hedir and B. Haddad, "Automatic system for radar echoes filtering based on textural features and artificial intelligence," *Meteorology and Atmospheric Physics*, vol. 129, no. 5, pp. 555–572, 2017.
[4] A. McGovern, K. L. Elmore, I. Gagne, David John, S. E. Haupt, C. D. Karstens, R. Lagerquist, T. Smith, and J. K. Williams, "Using Artificial Intelligence to Improve Real-Time Decision-Making for High-Impact Weather," *Bulletin of the American Meteorological Society*, vol. 98, no. 10, pp. 2073–2090, 2017.
[5] S. Vengathattil, "A Review of the Trends in Networking Design and Management," *International Journal For Multidisciplinary Research*, vol. 2, no. 3, p. 37456, 2020.

[6] U.-O. Dorj, K.-K. Lee, C. Jae-Young, and M. Lee, "The skin cancer classification using deep convolutional neural network," *Multimedia Tools and Applications*, vol. 77, no. 8, pp. 9909–9924, 2018.

[7] R. G. Smith and J. Eckroth, "Building AI Applications: Yesterday, Today, and Tomorrow," *AI Magazine*, vol. 38, no. 1, pp. 6–22, 2017.

[8] V. Rus, Z. Markov, and I. Russell, "Report on the 30th International Florida Artificial Intelligence Research Society Conference (FLAIRS-30)," *AI Magazine*, vol. 38, no. 3, pp. 70–71, 2017.

[9] M. Anderson, R. Barták, J. S. Brownstein, D. L. Buckeridge, H. Eldardiry, C. Geib, M. Gini, A. Isaksen, S. Keren, R. Laddaga, V. Lisy, R. Martin, D. Martinez, M. Michalowski, L. Michael, R. Mirsky, T. Nguyen, M. J. Paul, E. Pontelli, S. Sanner, A. Shaban-Nejad, A. Sinha, S. Sohrabi, K. Sricharan, B. Srivastava, M. Stefik, W. Streilein, N. Sturtevant, K. Talamadupula, M. Thielscher, J. Togelius, S. C. Tran, L. Tran-Thanh, N. Wagner, B. C. Wallace, S. Wilk, and J. Zhu, "Reports of the Workshops of the 31st AAAI Conference on Artificial Intelligence," *AI Magazine*, vol. 38, no. 3, pp. 72–82, 2017.

[10] J. Hernández-Orallo, M. Baroni, J. Bieger, N. Chmait, D. L. Dowe, K. Hofmann, F. Martínez-Plumed, C. Strannegård, and R. Thórissons, Kristinn, "A New AI Evaluation Cosmos: Ready to Play the Game?" *AI Magazine*, vol. 38, no. 3, pp. 66–69, 2017.

[11] T. G. Dietterich, "AAAI Presidential Address: Steps Toward Robust Artificial Intelligence," *AI Magazine*, vol. 38, no. 3, pp. 3–24, 2017.

[12] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and Brain Sciences*, vol. 40, 2017.

[13] J. E. Laird, C. Lebiere, and P. S. Rosenbloom, "A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics," *AI Magazine*, vol. 38, no. 4, pp. 13–26, 2017.

[14] K. D. Forbus and T. Hinrichs, "Analogy and Qualitative Representations in the Companion Cognitive Architecture," *AI Magazine*, vol. 38, no. 4, pp. 34–42, 2017.

[15] P. F. Bello and W. Bridewell, "There Is No Agency Without Attention," *AI Magazine*, vol. 38, no. 4, pp. 27–33, 2017.

[16] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis, "Mastering the game of Go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359,359A–359L, 2017.

[17] F. Chesani, P. Mello, and M. Milano, "Solving Mathematical Puzzles: A Challenging Competition for AI," *AI Magazine*, vol. 38, no. 3, pp. 83–96, 2017.

[18] E. Davis, L. Morgenstern, and C. L. Ortiz Jr, "The First Winograd Schema Challenge at IJCAI-16," *AI Magazine*, vol. 38, no. 3, pp. 97–98, 2017.

[19] L. Kotthoff, B. Hurley, and B. O'Sullivan, "The ICON Challenge on Algorithm Selection," *AI Magazine*, vol. 38, no. 2, pp. 91–93, 2017.

[20] S. Ganzfried, "Reflections on the First Man Versus Machine No-Limit Texas Hold 'em Competition," *AI Magazine*, vol. 38, no. 2, pp. 77–85, 2017.

[21] S. N. Minton, "The Value of AI Tools: Some Lessons Learned," *AI Magazine*, vol. 38, no. 3, pp. 99–101, 2017.

[22] Y. Wang and J. Zhang, "Exploring topics related to data mining on Wikipedia," *The Electronic Library*, vol. 35, no. 4, pp. 667–688, 2017.

[23] M. A. A. Rad and M. S. A. Rad, "Comparison of artificial neural network and coupled simulated annealing based least square support vector regression models for prediction of compressive strength of high-performance concrete," *Scientia Iranica.Transaction A, Civil Engineering*, vol. 24, no. 2, pp. 487–496, 2017.

[24] I. Peško, V. Mučenski, M. Šešlija, N. Radović, A. Vujkov, D. Bibić, and M. Krklješ, "Estimation of Costs and Durations of Construction of Urban Roads Using ANN and SVM," *Complexity*, vol. 2017, p. 13, 2017.

[25] D.-M. Petrosanu and A. PÎrjan, "IMPLEMENTATION SOLUTIONS FOR DEEP LEARNING NEURAL NETWORKS TARGETING VARIOUS APPLICATION FIELDS," *Journal of Information Systems & Operations Management*, pp. 155–169, 2017.

[26] G. Ceribasi, E. Dogan, U. Akkaya, and U. E. Kocamaz, "Application of trend analysis and artificial neural networks methods: The case of Sakarya River," *Scientia Iranica.Transaction A, Civil Engineering*, vol. 24, no. 3, pp. 993–999, 2017.

[27] B. Brożek and M. Jakubiec, "On the legal responsibility of autonomous machines," *Artificial Intelligence and Law*, vol. 25, no. 3, pp. 293–304, 2017.

[28] H. Prakken, "On the problem of making autonomous vehicles conform to traffic law," *Artificial Intelligence and Law*, vol. 25, no. 3, pp. 341–363, 2017.

[29] S. M. Solaiman, "Legal personality of robots, corporations, idols and chimpanzees: a quest for legitimacy," *Artificial Intelligence and Law*, vol. 25, no. 2, pp. 155–179, 2017.

[30] S. Vengathattil, "Ethical Artificial Intelligence - Does it exist?" *International Journal For Multidisciplinary Research*, vol. 1, no. 3, p. 37443, 2019.

[31] J. Kingston, "Using artificial intelligence to support compliance with the general data protection regulation," *Artificial Intelligence and Law*, vol. 25, no. 4, pp. 429–443, 2017.

[32] W. A. Kaal and E. P. M. Vermeulen, "HOW TO REGULATE DISRUPTIVE INNOVATION-FROM FACTS TO DATA," *Jurimetrics*, vol. 57, no. 2, pp. 169–209, 2017.

[33] I. Rahwan, "Society-in-the-loop: programming the algorithmic social contract," *Ethics and Information Technology*, vol. 20, no. 1, pp. 5–14, 2018.

[34] S. D. Baum, "On the promotion of safe and socially beneficial artificial intelligence," *AI & Society*, vol. 32, no. 4, pp. 543–551, 2017.

[35] T. Arnold and M. Scheutz, "The "big red button" is too late: an alternative model for the ethical evaluation of AI systems," *Ethics and Information Technology*, vol. 20, no. 1, pp. 59–69, 2018.

[36] P. Vamplew, R. Dazeley, C. Foale, S. Firmin, and J. Mummery, "Human-aligned artificial intelligence is a multiobjective problem," *Ethics and Information Technology*, vol. 20, no. 1, pp. 27–40, 2018.

[37] D. G. Johnson and M. Verdicchio, "AI, agency and responsibility: the VW fraud case and beyond," *AI & Society*, vol. 34, no. 3, pp. 639–647, 2019.

[38] E. Anthes, "THE SHAPE OF WORK TO COME," *Nature*, vol. 550, no. 7676, pp. 316–319, 2017.

[39] O. Churchill, "CHINA'S AI DREAMS," *Nature, suppl.nature jobs CAREER GUIDE CHINA*, vol. 553, no. 7688, pp. S10–S12, 2018.

[40] T. Walsh, "The Singularity May Never Be Near," *AI Magazine*, vol. 38, no. 3, pp. 58–62, 2017.

[41] I. Sitnicki, "Why AI shall emerge in the one of possible worlds?" *AI & Society*, vol. 34, no. 2, pp. 365–371, 2019.

[42] V. Galanos, "Singularitarianism and schizophrenia," *AI & Society*, vol. 32, no. 4, pp. 573–590, 2017.

[43] M. Wheeler, "The reappearing tool: transparency, smart technology, and the extended mind," *AI & Society*, vol. 34, no. 4, pp. 857–866, 2019.

[44] M. B. Fazi, "Can a machine think (anything new)? Automation beyond simulation," *AI & Society*, vol. 34, no. 4, pp. 813–824, 2019.

[45] V. Havlík, "The naturalness of artificial intelligence from the evolutionary perspective," *AI & Society*, vol. 34, no. 4, pp. 889–898, 2019.

[46] C. L. Oguego, J. C. Augusto, A. Muñoz, and M. Springett, "A survey on managing users' preferences in ambient intelligence," *Universal Access in the Information Society*, vol. 17, no. 1, pp. 97–114, 2018.

[47] S. Kanagarajan and S. Ramakrishnan, "Ubiquitous and Ambient Intelligence Assisted Learning Environment Infrastructures Development - a review," *Education and Information Technologies*, vol. 23, no. 1, pp. 569–598, 2018.

[48] G. Kessler, "Technology and the future of language teaching," *Foreign Language Annals*, vol. 51, no. 1, pp. 205–218, 2018.

[49] K. J. Luterbach, W. Xiao, H. Tang, and H. Chen, "Reflections on World Education Day 2017," *TechTrends*, vol. 62, no. 2, pp. 143–145, 2018.

[50] T. Bench-capon, "HYPO'S legacy: introduction to the virtual special issue," *Artificial Intelligence and Law*, vol. 25, no. 2, pp. 205–250, 2017.

[51] D. Walton and T. F. Gordon, "How Computational Tools Can Help Rhetoric and Informal Logic with Argument Invention," *Argumentation*, vol. 33, no. 2, pp. 269–295, 2019.

[52] A. Mardani, M. Nilashi, J. Antucheviciene, M. Tavana, R. Bausys, and I. Othman, "Recent Fuzzy Generalisations of Rough Sets Theory: A Systematic Review and Methodological Critique of the Literature," *Complexity*, vol. 2017, p. 33, 2017.

[53] J. Gonzalez-Cava, J. A. Reboso, J. Casteleiro-Roca, J. Calvo-Rolle, and J. A. Méndez Pérez, "A Novel Fuzzy Algorithm to Introduce New Variables in the Drug Supply Decision-Making Process in Medicine," *Complexity*, vol. 2018, p. 15, 2018.