# Synthesis of the Dissertation

*Toward Smaller, Lighter, and More Transparent AI*

**Context and Motivation.** Artificial Intelligence has rapidly transitioned from million-parameter CNNs to foundation models with billions–trillions of parameters. This growth delivers impressive capabilities yet strains compute, energy, and safety: models are expensive to train and deploy, hard to fit on edge devices, and difficult to interpret or steer. This dissertation tackles these pressures along three axes: making models *smaller* (via pruning with a new neuron design), *lighter* (via event-based, energy-efficient vision pipelines), and *more transparent* (via concept-based models enhanced through variational inference).

**I. Smaller: MAM Neurons for Aggressive Pruning.** I introduce the *Multiply-and-Max/min (MAM)* neuron, a map–reduce alternative to standard multiply-and-accumulate (MAC). MAM multiplies inputs by weights and reduces via the sum of the maximum and minimum contributions, concentrating computation on a few decisive connections. A *vanishing-contributions* training schedule smoothly interpolates from MAC to MAM, preserving trainability and accuracy. Across fully-connected layers of canonical computer-vision models (AlexNet, VGG-16) and ViT-B/16, non-structured magnitude and gradient-based pruning removes the vast majority of weights while maintaining competitive performance. Theoretical analysis establishes universal approximation for networks built with MAM layers.

**II. Lighter: Event-Based Vision with Dynamical Attention.** To reduce sensing and compute costs, I leverage neuromorphic *event cameras.* The work contributes a large, manually annotated person-detection dataset (PEDRo) and proposes *MESA* (Memory of Events through Spatial Attention), a dynamical attention pre-processing that forms decaying memory tensors and spatially modulated forgetting masks. MESA enables non-recurrent backbones to process sparse asynchronous events effectively, improving accuracy with modest memory/compute overhead and boosting downstream metrics (classification and detection) compared with recurrent or frame-accumulation baselines.

**III. More Transparent: Variational Concept-Based Models.** I propose *V-CEM*, a variational extension of Concept Embedding Models that learns dense, well-clustered concept representations via approximate posteriors matched to a prior. V-CEM attains in-distribution accuracy competitive with black-box models on several vision and language tasks, improves robustness under noise and distribution shifts, and *enhances intervenability*: human edits to concept predictions reliably propagate to correct model outputs. New evaluation criteria (e.g., concept cohesiveness) characterize when such interventions are effective.

**Conclusions and Outlook.** This dissertation shows that pruning-centered architectures (via MAM) can retain high accuracy after extreme sparsification, opening paths to ultra-compact sub-networks in larger systems. Event-based pipelines (PEDRo + MESA) demonstrate that asynchronous sensors and dynamical attention can deliver strong accuracy at low cost, encouraging broader adoption in embedded perception. Finally, V-CEM provides a principled route to transparent, steerable AI, combining black-box-level accuracy with reliable concept-level intervention. Future work includes: hardware co-design for MAM layers; richer event representations and real-time neuromorphic processing; multi-modal V-CEM with generative decoders; and explicit modeling of concept dependencies to further enhance accountability and trustworthiness of AI systems.

**Keywords:** pruning, universal approximation, event-based cameras, MESA, concept-based explainability.