Appendix A. AI Index

Table 10: AI Technology Sector Index

Ticker	Company Name	Primary Business Segment
NVDA	NVIDIA Corporation	GPU/AI accelerators, data center infrastructure
MSFT	Microsoft Corporation	Cloud computing (Azure), generative AI (Copilot)
GOOGL	Alphabet Inc.	Search, advertising, cloud (Google Cloud), LLMs (Gemini)
AMZN	Amazon.com Inc.	Cloud computing (AWS), e-commerce, AI services
META	Meta Platforms Inc.	Social media, advertising, LLM research (Llama)
AAPL	Apple Inc.	Consumer electronics, mobile devices, AI chips
INTC	Intel Corporation	Semiconductor manufacturing, processors, AI chips
AMD	Advanced Micro Devices	Semiconductors, GPU/AI accelerators
QCOM	Qualcomm Inc.	Mobile processors, automotive semiconductors
ORCL	Oracle Corporation	Cloud infrastructure, database services
IBM	International Business Machines	Enterprise IT, quantum computing research
CRM	Salesforce Inc.	Enterprise CRM, AI-powered analytics
ADBE	Adobe Inc.	Creative software, generative AI (Firefly)
NOW	ServiceNow Inc.	Enterprise workflow automation, AI capabilities
SNOW	Snowflake Inc.	Cloud data platform, data analytics
INTU	Intuit Inc.	Financial software, AI-driven automation
AVGO	Broadcom Inc.	Networking semiconductors, data center infrastructure
ANET	Arista Networks	Cloud networking, data center solutions
APP	AppLovin Corporation	Mobile application analytics, AI advertising
PLTR	Palantir Technologies	Data analytics, AI/ML for enterprise
AI	C3.ai Inc.	AI/ML infrastructure and analytics
SMCI	Super Micro Computer	AI server hardware, data center systems
ZM	Zoom Video Communications	Video conferencing, AI-enhanced communication
CSCO	Cisco Systems	Networking infrastructure, cybersecurity
DELL	Dell Technologies	Enterprise IT hardware, servers
HPE	Hewlett Packard Enterprise	Enterprise servers, data center, infrastructure
SOUN	SoundHound AI	Conversational AI, voice recognition

Appendix B. Mapping of Literature and Research Directions

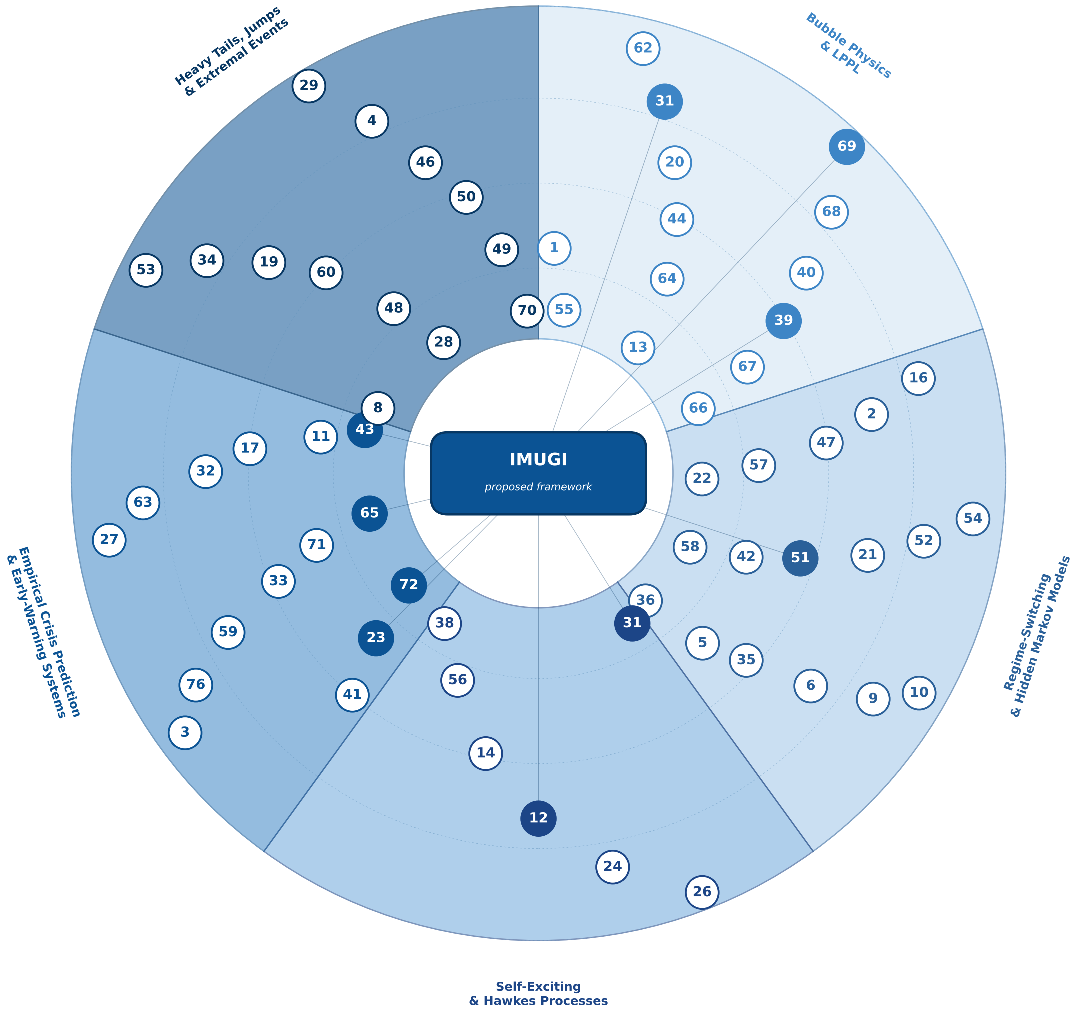

Figure 5: Mapping of the most relevant references and research directions. The proposed IMUGI framework integrates contributions from five research streams: bubble physics and log-periodic power law (LPPL) modelling; regime-switching and hidden Markov models; self-exciting and Hawkes processes; empirical crisis prediction and early-warning systems; and heavy-tailed, jump, and extremal-event analysis. Filled markers denote references the framework directly extends or operationalises; lines from the centre node indicate direct architectural integration.

Appendix C. Computational Procedure

This appendix specifies implementation details required for replication that complement the model definitions in Section 3: convergence rules, numerical safeguards, damping coefficients, and box constraints.

SOHMM: EM Convergence Criteria

Algorithm 1 SOHMM EM Fitting

Require: $\theta^{(0)}$, $\tau = 10^{-4}$, $I_{\max} = 30$

Ensure: $\hat{\theta}$

$i \leftarrow 0$, $\theta^\star \leftarrow \theta^{(0)}$, $\ell^\star \leftarrow -\infty$

repeat

    E-step: $q_t(j,k) \leftarrow p(z_{t-1} = j, z_t = k \mid X, \theta^{(i)})$

    $\ell^{(i)} \leftarrow \log p(X \mid \theta^{(i)})$

    if $i > 3$ **and** ($\ell^{(i)} < \ell^{(i-1)} - 100$ **or** $\neg$finite($\ell^{(i)}$)) **then**

        $\theta^{(i)} \leftarrow \theta^\star$; **break** ▷ degradation guard

    end if

    if $i > 3$ **and** $|\ell^{(i)} - \ell^{(i-1)}| < \tau$ **then**

        break ▷ convergence

    end if

    if $\ell^{(i)} > \ell^\star$ **then**

        $\ell^\star \leftarrow \ell^{(i)}$, $\theta^\star \leftarrow \theta^{(i)}$

    end if

    M-step: $\theta^{(i+1)} \leftarrow \arg\max_\theta \mathbb{E}_q[\log p(X, Z \mid \theta)]$

    $i \leftarrow i + 1$

until $i = I_{\max}$

Numerical safeguards. The multivariate-t log-density is clamped to $[-700, 700]$. Cholesky factorisation uses an adaptive ridge $\Sigma_k + \delta I_D$ with $\delta \in \{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}$, attempted in increasing order, with a diagonal-only fallback. Probability vectors are clipped to $[10^{-12}, 1]$ before normalisation; Ψ_k to $[10^{-4}, 1 - 10^{-4}]$. The M-step covariance ridge is

$\lambda = 10^{-2}$.

SOHMM: Online Damping Coefficients

Algorithm 2 SOHMM Online Parameter Update

Require: posterior γ_t, KL surprise $\mathcal{S}_t$, sufficient statistics (S_w, S_m, S_x)

Ensure: adapted $(\mu_k, \Psi_k, \Sigma_k)$

1: $\mathcal{S}_t \leftarrow \mathrm{clip}(\mathcal{S}_t,\, 0,\, 10)$

2: $\eta_t \leftarrow \eta_{\text{base}} + (\eta_{\max} - \eta_{\text{base}})\, \sigma(5(\mathcal{S}_t - 0.3)), \quad$ with $(\eta_{\text{base}}, \eta_{\max}) = (0.002, 0.20)$

3: $\hat{\mu}_k \leftarrow S_x/S_m, \quad \hat{\Psi}_k \leftarrow \mathrm{clip}(S_m/S_w,\, 10^{-4},\, 1 - 10^{-4})$

4: $\mu_k \leftarrow (1 - \kappa_\mu \eta_t)\, \mu_k + \kappa_\mu \eta_t\, \hat{\mu}_k$

5: $\Psi_k \leftarrow (1 - \kappa_\Psi \eta_t)\, \Psi_k + \kappa_\Psi \eta_t\, \hat{\Psi}_k$

6: $\Sigma_k \leftarrow (1 - \kappa_\Sigma \eta_t)\, \mathrm{diag}(\Sigma_k) + \kappa_\Sigma \eta_t\, \mathrm{diag}\big((x_t - \mu_k)^2 \odot m_t\big) + \lambda I_D$

The damping coefficients are $(\kappa_\mu, \kappa_\Psi, \kappa_\Sigma) = (0.10, 0.05, 0.05)$. They multiply η_t rather than replace it, so high-surprise observations produce fast parameter adjustment while low-surprise observations are doubly damped. The covariance update is restricted to the diagonal: off-diagonal entries are inherited from the batch-fitted estimates and not refined online, keeping per-step cost at $\mathcal{O}(KD)$ rather than $\mathcal{O}(KD^2)$.

Q-Hawkes: Box Constraints and Initial Values

Algorithm 3 Q-Hawkes Estimation

Require: event times $\{t_i\}_{i=1}^N$, horizon T, $\lambda_{\text{reg}} = 10^{-5}$

Ensure: $\hat{\theta}$

1: $\hat{\mu}_0 \leftarrow 0.5\, N_{\text{events}}/T$

2: $\hat{\alpha}_0 \leftarrow 0.25, \quad \hat{\beta}_0^{\text{diag}} \leftarrow 0.05, \quad \hat{\beta}_0^{\text{rank1}} \leftarrow 0.025$

3: $\hat{\theta}_0^{(0)} \leftarrow 0.25, \quad \hat{\theta}_0^{(1)} \leftarrow 0.30, \quad \hat{\theta}_0^{(2)} \leftarrow 0.20$

4: Box constraints: $\mu \geq 10^{-6}$, $\alpha, \beta^{\text{diag}}, \beta^{\text{rank1}} \geq 0$, $\theta^{(0)}, \theta^{(1)}, \theta^{(2)} \in [10^{-3}, 10]$

5: Solve with L-BFGS-B; stop at scipy's default convergence criterion ($\|\nabla\mathcal{L}\|_{\text{proj}} < 10^{-5}$) or 1,000 iterations

Computational Cost

The SOHMM forward-backward recursion is $\mathcal{O}(TK^2)$ in time and memory, the M-step covariance update $\mathcal{O}(KTD^2)$, and the Q-Hawkes likelihood evaluation $\mathcal{O}(N_{\text{events}}^2 + N_{\text{quad}}N_{\text{events}})$ per L-BFGS-B call. The meta-learner is negligible. End-to-end cold-start training of the three layers completes in approximately 60–90 seconds, with per-day online inference in the 10–40 ms range. LPPL features are precomputed once on the full sample using strictly causal windows; this stage, dominated by the differential-evolution optimisation, is the principal offline cost.