ACM Transactions on Information Systems
Special Section on Efficiency in Neural Information Retrieval

Guest Editors:
- **Dr. Sebastian Bruch**, Pinecone, United States of America, sbruch@acm.org
- **Prof. Claudio Lucchese**, Ca' Foscari University of Venice, Italy, claudio.lucchese@unive.it
- **Prof. Maria Maistro,** University of Copenhagen, Denmark, mm@di.ku.dk
- **Dr. Franco Maria Nardini**, ISTI-CNR, Italy, francomaria.nardini@isti.cnr.it

The rise of deep neural networks and self-supervised learning in recent years have brought about a paradigm shift in Information Retrieval. From retrieval to ranking, question answering to recommendation, search to conversational agents, models trained on hand-crafted features have given way to complex neural networks built of millions of parameters that are capable of learning granular features from raw data. While this transition has led to large gains in the efficacy in various tasks, it has been done so often at the expense of training and inference efficiency. With deep models embedded in evermore applications and devices with a drive towards ever higher efficacy, the rise in costs has a tangible though often underreported impact on researchers, practitioners, users, and more importantly, the environment. It is therefore unsurprising that the difficult balancing act between celebrating effectiveness and seeking efficiency has resurrected old research questions from the field with a renewed urgency.

The aim of this Special Section is to engage with researchers in Information Retrieval, Natural Language Processing, and related areas and gather insight into the core challenges in measuring, reporting, and optimizing all facets of efficiency in Neural Information Retrieval (NIR) systems, including time-, space-, resource-, sample- and energy-efficiency, among other factors. While researchers in the field have assiduously explored the Pareto frontier in quality and efficiency in other contexts for decades, we believe that the neural dimension introduces new hurdles. This special section solicits perspectives from active researchers to advance our understanding of and to overcome efficiency challenges in NIR. In particular, researchers are encouraged to examine the ever-growing model complexity through appropriate empirical analysis, to propose models that require less data, computational resources, and energy for training and fine-tuning with similarly efficient inference, to ask if there are meaningful simplifications of the existing training processes or model architectures that lead to comparable quality, and explore a multi-faceted evaluation of NIR models from quality to all dimensions of efficiency with standardized metrics.

Topics
We welcome submissions on the following topics, including but not limited to:
- Novel NIR models that reach competitive quality but are designed to provide efficient training or inference;
- Efficient NIR models for decentralized IR tasks such as conversational search;
- Efficient NIR models for IR-related tasks such as question answering and recommender systems;
- Efficient NIR for resource-constrained devices;
- Scalability of NIR systems;
- Efficient NIR for text and cross-modal search;
- Strategies to optimize training or inference of existing NIR models;
- Sample-efficient training of NIR models;
- Efficiency-driven distillation, pruning, quantization, retraining, and transfer learning;

- Empirical investigation of the complexity of existing NIR models through an analysis of quality, interpretability, robustness, and environmental impact; and,
- Evaluation protocols for efficiency in NIR.

Important Dates
- Open for Submissions: Aug 1, 2022
- Submissions deadline: Feb 28, 2023
- First-round review decisions: May 31, 2023
- Deadline for minor revision submissions: June 30, 2023
- Deadline for major revision submissions: Aug 31, 2023
- Notification of final decisions: Sept 30, 2023
- Tentative publication: 2024

Submission Information

Authors can submit their manuscripts via https://dl.acm.org/journal/tois. Submissions to this special issue will follow the regular TOIS submission guidelines (https://dl.acm.org/journal/tois/author-guidelines). Submissions must be accompanied by a cover letter containing all of the following: (1) Confirm that the paper is not currently under submission at another journal or conference. (2) Confirm that the paper is substantially different from any previously published work. (3) Confirm that none of the co-authors is a Guest Editor for this special session. (4) Disclose possible conflicts of interest with Guests Editors. The review process will be single-blind. Strict policies will be followed for plagiarism, submission confidentiality, reviewer anonymity, prior and concurrent paper submission based on the guidelines.

Papers with a "Major Revision" decision should be resubmitted within three months, and with a "Minor Revision" decision should be resubmitted within one month. Revised submissions must be accompanied with a detailed response to reviewers explaining what revisions were implemented. The editors will conduct a second-round review process and give the decision (accept or reject or need further revision) in one month.

For questions and further information, please contact Dr. **Sebastian Bruch** at **sbruch[at]acm[dot]org**.