Journal of Edge Computing, 2026, Vol. **5**, Iss. 2, pp. 189–214 10.55056/jec.1518

doors 2026: a quantitative review of the 6th Edge Computing Workshop, with JEC Volume 5 Issue 2 (2026) highlights

Serhiy O. Semerikov[1,2,3,4,5], Tetiana A. Vakaliuk[3,2,1,5]

[1]*Kryvyi Rih State Pedagogical University, 54 Universytetskyi Ave., Kryvyi Rih, 50086, Ukraine*
[2]*Institute for Digitalisation of Education of the NAES of Ukraine, 9 M. Berlynskoho Str., Kyiv, 04060, Ukraine*
[3]*Zhytomyr Polytechnic State University, 103 Chudnivsyka Str., Zhytomyr, 10005, Ukraine*
[4]*Center for Information-analytical and Technical Support of Nuclear Power Facilities Monitoring of the NAS of Ukraine, 34a Palladin Ave., Kyiv, 03142, Ukraine*
[5]*Academy of Cognitive and Natural Sciences, 54 Gagarin Ave., Kryvyi Rih, 50086, Ukraine*

Abstract. This article reviews the 6th Edge Computing Workshop (doors 2026), held online on 3 April 2026, whose seven papers were published in the ACM International Conference Proceedings Series. It is a substantially extended and revised version of the workshop report in the proceedings. We combine the review record, a recording of the whole session, a facet analysis of the papers by an ensemble of six language-model coders with reported reliability, a coded analysis of the 23 questions asked in the discussion, and a positioning of each paper against 51 recent peer-reviewed works. doors 2026 accepted 7 of 16 submissions (44 %), the lowest rate among the editions with known submission counts. Most papers place computation on the end device, all were evaluated in the laboratory, and one links its code. Three papers address combinations without a direct precedent in recent work, but the evaluation of five is weaker than the norm of their themes. The discussion took 33 % of the time within the rounds; all questions came from the floor, concentrated on measurement, operating conditions and the placement of computation, and several anticipated limitations stated in the published papers. We derive a research agenda of seven items centred on testing under deployment conditions and on making explicit where each part of an edge system runs. The article closes with an overview of the papers in Volume 5 Issue 2 (2026) of the *Journal of Edge Computing*.

Keywords: edge computing, workshop review, doors 2026, discourse analysis, conference discussion, edge AI, fog computing, Internet of Things

1. Introduction

Edge computing moves computation from distant data centres towards the devices that produce data, trading the elasticity of the cloud for lower latency, less traffic and more local control [64, 74]. Between the device and the cloud, fog nodes take on part of the work [13]; at the network edge, servers co-located with mobile infrastructure do the same [50]; and the arrival of machine learning on all these tiers has turned edge computing into edge intelligence [92]. Research on these questions is spread across hardware, networking, systems and machine learning, which makes small specialised meetings useful places to see where a community stands.

The Edge Computing Workshop, doors, is one such meeting. Since 2021 it has met every spring, hosted from Zhytomyr, and its proceedings have appeared in CEUR Workshop Proceedings [65, 84, 86, 87]. The sixth edition, doors 2026, met online on 3 April 2026, and its proceedings were the first

This article is a substantially extended and revised version of the workshop report published in the front matter of the doors 2026 proceedings [88].

0000-0003-0789-0272 (S. O. Semerikov); 0000-0001-6825-4697 (T. A. Vakaliuk)
semerikov@gmail.com (S. O. Semerikov); tetianavakaliuk@acnsci.org (T. A. Vakaliuk)
https://acnsci.org/semerikov (S. O. Semerikov); https://acnsci.org/vakaliuk (T. A. Vakaliuk)

Received 2025-10-25 | Accepted 2026-11-20 | Published 2026-11-21 | Version of record 2026-11-21

© Copyright for this article by its authors, published by the Academy of Cognitive and Natural Sciences. This is an Open Access article distributed under the terms of the Creative Commons License Attribution 4.0 International (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

of the series to appear in the ACM International Conference Proceedings Series [88]. Each edition has been followed by a report in this journal [68, 83, 85]. These reports describe the call, the review and the accepted papers, and one combined such an overview with a bibliometric analysis of the field [79]; none examined what happened in the sessions or measured the papers against the wider literature.

This article does both for doors 2026. It is a substantially extended and revised version of the workshop report in the front matter of the proceedings [88], whose text it does not reuse. Like our recent surveys of edge computing research [66, 70], it treats its subject as data and asks three questions:

RQ1. What problems do the doors 2026 papers address, at which tier of the edge continuum, with what kind of evidence, and how do they compare with work published elsewhere in 2024–2026?

RQ2. What did the discussion at the workshop add: what was asked, by whom, how were the questions answered, and how does the discussion relate to the published papers?

RQ3. How does doors 2026 compare with the earlier editions, and what research agenda follows from its papers, its discussion and the literature together?

The answers rest on the review record, on a recording of the full session and on a coded analysis of the papers and the discussion, with the coding done by an ensemble of independently trained language models and its reliability reported. The article contributes a quantitative profile of the workshop and its series (section 3), a facet analysis of its papers (section 4), a summary of each paper with what the session added (section 5), an analysis of the discussion (section 6), a positioning of each paper against recent work (section 7) and a research agenda (section 8). Section 2 describes the methods, section 9 their limits, and section 10 answers the research questions. As the first article of the journal's Volume 5 Issue 2, it closes with an overview of that issue (section 11).

2. Methodology

The study combines five sources: the published papers, the review record, a recording of the session, the workshop report, and recent literature on the papers' themes. Every number in this article is computed by a script from these sources and written into the text as a macro, and the scripts check that related numbers agree, for instance that accepted and rejected papers add up to the reviewed ones. Scripts, codebooks and derived data are available as described in the data availability statement.

2.1. Corpus and review record

The corpus consists of the seven papers published in the proceedings [6, 40, 41, 48, 49, 75, 80], all under a CC BY 4.0 licence. We worked from their camera-ready LaTeX sources, which give exact titles, authors, affiliations, cited references and figure counts; page and word counts come from the published PDF files. Cited references were parsed from each paper's bibliography and classified by type and year, and links to code or data repositories were detected by pattern. The workshop report in the front matter of the proceedings [88] supplied the call, the committees and the programme; its text is not reused here.

The review record is the export of the submission system. We report it only in aggregate: no review text is quoted, and reviewers are replaced by pseudonyms before any processing. A paper submitted twice is counted once. The export does not distinguish papers rejected without review from those withdrawn, so both are reported together as not sent to review.

2.2. Facet coding of the papers

Each paper was coded on the eight facets of table 4, whose codebook was fixed before any coding began. To reduce dependence on a single judge, the coding was done by five open-weight language models running on a local server (`gemma4:26b`, `granite4.1:8b`, `llama3.3:70b`, `nemotron3:33b`, `qwen3.8-flash-next:125b-a6b-q4_K_M`) and by a separate pass of Claude, the assistant used in preparing this study. The local models received the codebook and a dossier of the paper's text extracted from its PDF. They ran one at a time, at temperature 0 with a fixed seed, and had to reply in JSON with valid codes; an invalid reply was resampled with a new seed up to three times. The consensus code of each cell is the majority code, with ties broken by the Claude pass. We report Krippendorff's α for nominal data [36, 45] and the mean pairwise Gwet's AC1 [35], which is less sensitive to skewed code distributions [21, 28], each with a percentile bootstrap interval over 1,000 resamples [25]. Cells chosen by three or fewer of the six coders were adjudicated against the full text. With seven papers, these coefficients measure how far independently built models agree on the codebook; they are not a substitute for human inter-rater reliability, and no claim in this article rests on them alone. Language models have been shown to approach or exceed crowd annotators on comparable annotation tasks [32, 93], but not to replace expert coders.

2.3. Recording, transcription and segmentation

The session was held on Zoom and recorded in full (2:21:15). We transcribed the audio twice. The primary transcript was produced with Whisper [60] (model `small.en`) with word-level timestamps; a second reading, used to check and repair the primary one, was produced with the `large-v3` model through the faster-whisper implementation. An earlier automatic transcript without timestamps served only for orientation. Neither transcript carries speaker labels, so speakers were identified from the chairs' invitations and from the course of the conversation; where this was not possible, an utterance is marked as unattributed.

The session was divided into rounds, one per paper, and each round into a handover (the chair's announcement and any screen-sharing or audio problems), the talk, and the discussion. Boundaries were placed by reading the transcript around recurring formulas, such as a chair inviting the next speaker or a presenter thanking the audience for their attention, and every boundary is recorded with a fragment of its text, so that it can be checked against the recording. The delivered order agreed with the programme.

2.4. Coding of the discussion

The unit of analysis is an *ask*: one turn, or a run of consecutive turns by the same speaker other than the presenter, that requests information, challenges the work or recommends something. Procedural turns, such as calling on the next speaker or thanking, are not coded. Each ask was coded for its act, its topic and the response it received, following a code frame adapted from studies of peer review and conference discussion [34, 38, 42, 44, 59, 89]. The topic codes were recalibrated inductively for edge computing [14], and the response codes describe the content of an answer, never its fluency or the speaker's competence. The codebook was fixed before coding; one code, for an ask that received no answer, was added before any ask had been coded, because the original frame conflated a missing answer with a gap in the recording.

The asks were extracted by reading all 7 discussions in full and recorded with the transcript segments of the ask and of the response. The same panel of local models coded them blind: each model saw only the talk title and the automatic transcripts of the ask and the response, in random order and under opaque identifiers, never the asker's role. The Claude pass that extracted the asks also coded them; it was not blind to the speakers, and we report its agreement with the panel separately. Reliability is computed as for the facets, with the same caveat about small numbers.

Quotations carry the elapsed recording time and one of three tiers: V for verbatim, R for repaired, where recognition errors are corrected in square brackets after checking both transcripts, and P for

a paraphrase. Numbers and technical terms are never repaired; where one is doubtful, the ask is flagged and no claim rests on it. Presenters and participants appear by role only. Questions from the floor are reported in aggregate, and no answer is quoted to illustrate a weakness. The session was recorded with Zoom, which shows every participant that a recording is in progress, and the recording is used here only in the derived form described above. Before publication, everything this article says about each paper, including the questions put to its presenter, was offered to the paper's corresponding author for correction.

2.5. Positioning against recent literature

For each paper we defined a theme and searched for peer-reviewed work published from 2024 to 2026 on that theme, by topic only, mainly in OpenAlex and supplemented by other academic search services. We kept seven to eight works per theme, with at most one survey and with preprints admitted only where peer-reviewed work was scarce and labelled as such. Before any verdict was written, we fixed a two-part rule: the contribution is called distinctive, aligned or behind relative to the retrieved works, and the evaluation stronger, comparable or weaker than the evidence those works provide. Each call rests on at least two works. The search was carried out by a Claude agent under these rules; the verdicts are ours.

3. doors 2026 at a glance

3.1. The series, 2021–2026

The Edge Computing Workshop has met every spring since 2021 in Zhytomyr, hosted by Zhytomyr Polytechnic State University. Table 1 summarises its six editions. The first edition shared a proceedings volume with the Workshop on Quantum Information Technologies [65]. The war reduced the 2022 edition to a single accepted paper [83]. The editions of 2023 to 2025 appeared as volumes of CEUR Workshop Proceedings [84, 86, 87], and in 2026 the proceedings moved to the ACM International Conference Proceedings Series [88]. Across the six editions the workshop has accepted 46 papers. Submissions peaked in 2025, and the 2026 acceptance rate of 44 % is the lowest of the editions for which submission counts are known.

Table 1
The doors workshop series, 2021–2026. Accepted papers in 2025 are split into regular and short papers; 2021 counts the doors part of the joint QuaInT+doors volume; no proceedings volume was identified for 2022.

Year	Edition	Proceedings	Submitted	Accepted	Rate
2021	1st (joint QuaInT+doors)	CEUR-WS Vol-2850	14	7	50%
2022	2nd	JPCS Vol-2288	2	1	50%
2023	3rd	CEUR-WS Vol-3374	12	9	75%
2024	4th	CEUR-WS Vol-3666	19	9	47%
2025	5th	CEUR-WS Vol-3943	28	13 (9+4)	46%
2026	6th	ACM ICPS	16	7	44%

3.2. Submissions and review

The 2026 call asked for double-blind review by at least two members of the programme committee [88]. The submission system holds 17 entries, one of which is a second submission of the same paper, so 16 distinct papers were submitted. Of these, seven were not sent to review, nine were reviewed and seven were accepted: 44 % of all submissions and 78 % of the reviewed ones. The 28 reviews were written by eleven committee members, each of whom wrote between two and four reviews; every reviewed paper received three or four reviews, and authors replied to reviews seven times.

Reviewers scored eight criteria, an overall evaluation and their own expertise (table 2). Accepted papers scored higher than the papers rejected after review on nine of the nine criteria other than expertise; on the overall evaluation the means were 3.68 and 2.83. The median review ran to 96 words.

Table 2
Mean review scores per criterion (novelty on a 1–4 scale from 'published before' to 'surprisingly new'; all other criteria 1–5): the 7 accepted papers (22 reviews) and the 2 papers rejected after review (6 reviews).

Criterion	Accepted	Rejected after review
Relevance to the workshop	4.14	3.50
Content matches title	4.09	3.67
Abstract informativeness	3.82	3.33
Scientific novelty	2.50	2.00
Sources	3.41	3.33
Figures and tables	3.73	3.33
Scientific English	3.64	3.17
Formatting	3.95	3.17
Overall evaluation	3.68	2.83
Reviewer expertise	3.23	3.33

3.3. The programme as planned and as delivered

The workshop met online on 3 April 2026 in a single session. The programme allotted seven slots of 15 minutes, one per paper in proceedings order and without a break, from 14:10 to 15:55 local time. Figure 1 compares this plan with the session as recorded. The recording runs for 141 minutes; after 8 minutes of waiting, the chairs opened the session at 14:02. All papers were presented in the planned order. The first talk began five minutes ahead of schedule, and five of the seven talks started within five minutes of their planned time.

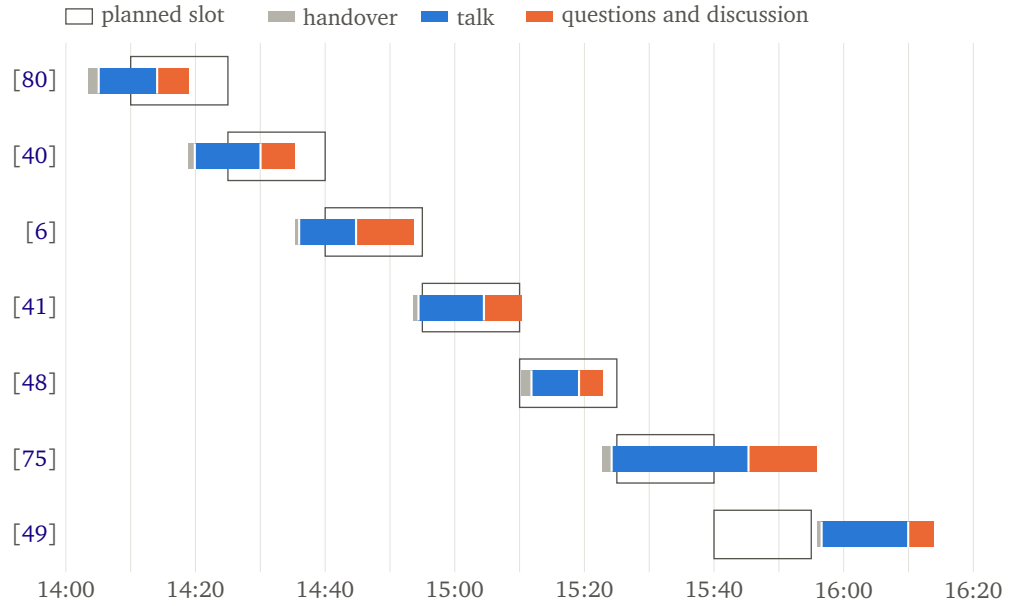


Figure 1: The session as planned and as delivered, in local time (UTC+3). Outlines mark the planned 15-minute slots; filled bars show the delivered handover, talk and discussion of each paper, timed from the recording. The chairs opened the session at 14:02.

Within the rounds, 80 minutes went to the talks (61 %), 43 to questions and discussion (33 %) and 8 to handovers (6 %), mostly spent on screen sharing and audio checks. Talks lasted between seven

and 21 minutes, discussions between four and eleven. Six of the seven rounds ran over their slot, all but one by at most three minutes. The exception was the round of [75], which took 33 minutes, 21 of them for the talk. It delayed the last talk by 17 minutes, and the session closed 19 minutes after the planned end. The recording shows no discussion closed for lack of time: the only remarks about time, late in the session, asked for a short question and answer, and several times the chairs asked whether anyone else had a question before moving on. Discussion therefore took a third of the time within the rounds (section 6).

3.4. Programme committee

The programme committee had 19 members from 13 countries, 15 of them working outside Ukraine; the two workshop chairs were among them:

- Olexander Barmak (Khmelnytskyi National University, Ukraine) [9]
- Ken Birman (Cornell University, USA) [76]
- Aleksandr Cariow (West Pomeranian University of Technology, Poland) [16]
- Josef Cernohorsky (Technical University of Liberec, Czech Republic) [19]
- Debashis De (Maulana Abul Kalam Azad University of Technology, India) [43]
- Balaji Shesharao Ingole (Data Flow, USA) [22]
- Valerii Kontsedailo (Inner Circle, Netherlands) [81]
- Vyacheslav Kryzhanivskyy (R&D Seco Tools AB, Sweden) [46]
- Nagender Kumar Suryadevara (University of Hyderabad, India) [78]
- Nadiia Lobanchykova (Zhytomyr Polytechnic State University, Ukraine) [58]
- Mykhailo Medvediev (ADA University, Azerbaijan) [82]
- Franco Milano (University of Florence, Italy) [55]
- Bongkyo Moon (Dongguk University, Republic of Korea) [5]
- Michael O'Grady (University College Dublin, Ireland) [91]
- Serhiy O. Semerikov (Kryvyi Rih State Pedagogical University, Ukraine) [71]
- Etibar Seyidzade (Azerbaijan Technical University, Azerbaijan) [29]
- Tetiana A. Vakaliuk (Zhytomyr Polytechnic State University, Ukraine) [67]
- Xianzhi Wang (University of Technology Sydney, Australia) [54]
- Alejandro Zunino (ISISTAN – UNCPBA & CONICET, Argentina) [37]

4. Corpus profile

4.1. Inventory, authorship and references

The proceedings contain seven papers on 66 pages, about 40,249 words in all (table 3). Their 27 authors come from 15 institutions in four countries (Germany, India, USA and Ukraine), and two papers are internationally co-authored; 26 authors give an ORCID. The papers cite 153 sources, between 14 and 33 per paper (median 24). The reference base is recent: the median cited work appeared in 2024, and 73 % of cited works were published in 2021 or later. Of the cited works, 84 are journal articles, 31 are conference papers and 38 are other sources such as books, preprints, standards and web pages; 68 % carry a DOI. Only one paper links its source code, and five papers include a declaration on the use of generative AI.

Table 3
The doors 2026 corpus: the seven proceedings papers in proceedings order.

Paper	Pages	Authors	Countries	Figures	Tables	Refs	Code
Vakaliuk, Andreiev et al. [80]	5	5	Ukraine	4	0	16	–
Kargin et al. [40]	8	4	Ukraine	5	3	27	–
Alekperov et al. [6]	10	5	Germany, Ukraine	10	9	24	–
Karnaukh et al. [41]	7	4	Ukraine	4	0	14	–
Lipianina-Honcharenko et al. [48]	10	4	Ukraine	3	5	14	–
Shvaika, Shvaika and Artemchuk [75]	13	3	USA, Ukraine	11	0	25	–
Mandal and Shende [49]	13	2	India	8	12	33	yes
Total	66	27	4	45	29	153	1

4.2. Facet profile

Each paper was coded on eight facets: where its computation runs, the constraint it mainly addresses, its primary technique, how it is evaluated and on what platform, which operating conditions it examines, whether its artefacts are available, and how mature it is. The procedure is described in section 2. Table 4 gives the consensus code for each paper and the number of the six coders who chose it.

Table 4
Facet profile of the seven papers: consensus of 6 coders (majority code; ties, 5 of 56 cells, broken by the Claude pass), with the number of coders supporting each code as a superscript and Krippendorff's α and mean pairwise Gwet's AC1 per facet. Codes supported by three or fewer coders were adjudicated against the full text; † marks the codes changed by adjudication. F8 is maturity on a 1–5 scale from concept to deployment. Full definitions are in the supplementary codebook.

Facet	[80]	[40]	[6]	[41]	[48]	[75]	[49]	α	AC1
F1 placement	device^{5}	multi-tier^{4}	multi-tier^{5}	device^{2}	device^{4}	cloud^{3}	device^{3}	0.18	0.23
F2 constraint	latency^{3}	uncertainty^{5}	latency^{6}	latency^{5}	few labels^{3}	eng. effort^{5}	throughput^{1†}	0.39	0.43
F3 technique	signal proc.^{6}	fuzzy/cogn.^{3}	architecture^{6}	analytical^{4}	frozen FM^{3}	LLM^{4}	runtime ctrl^{5}	0.53	0.52
F4 evaluation	controlled^{4}	demo/case^{4}	controlled^{4}	simulated^{6}	controlled^{4}	demo/case^{3}	controlled^{5}	0.32	0.38
F5 platform	target hw^{6}	target hw^{5}	target hw^{6}	simulator^{5}	desktop^{3}	none^{1†}	desktop^{4}	0.38	0.53
F6 conditions	nominal^{4}	nominal^{4}	varied^{6}	varied^{6}	varied^{5}	nominal^{3}	varied^{5}	0.34	0.51
F7 artefact	none^{5}	unlinked^{3}	unlinked^{4}	none^{4}	none^{4}	unlinked^{4}	linked^{5}	0.24	0.31
F8 maturity	3^{5}	3^{4}	3^{4}	2^{4}	3^{4}	3^{3}	3^{4}	0.04	0.25

Three features stand out. First, the programme is technically diverse: no two papers share a primary technique, which ranges from pixel-domain change detection to large language models. Second, most of the work sits at the device end of the continuum. Four papers place computation on the end device [41, 48, 49, 80], two divide it between tiers [6, 40], and one runs a cloud-hosted language model at design time while leaving execution to the IoT platform [75]. Latency is the constraint addressed most often [6, 41, 80]; the others are decision-making under asynchrony [40], scarce labelled data [48], engineering effort [75] and throughput [49].

Third, the evidence comes from the laboratory. Four papers compare alternatives on their own data or workloads [6, 48, 49, 80], two demonstrate a working system [40, 75], and one relies on network simulation [41]; none evaluates on a held-out benchmark with standard metrics. Three papers measure on the target edge hardware [6, 40, 80], two on a desktop or server standing in for it [48, 49], and one reports no measurements [75]. Operating conditions are varied in controlled tests in four papers [6, 41, 48, 49], and no paper reports field or deployment conditions. The coders placed six papers at level 3 of the maturity scale, a prototype validated on offline data or in the laboratory, and one at level 2, a proof of principle by analysis and simulation [41]. Only one paper links a public repository for its own contribution [49], in line with the inventory in section 4.1.

The coders agreed only moderately. Krippendorff's α ranged from 0.04 for maturity to 0.53 for technique, with a mean of 0.30; the mean pairwise Gwet's AC1 was 0.40. No facet reaches 0.667, the lowest value Krippendorff accepts for tentative conclusions [45], and with seven papers the bootstrap intervals [25] are wide: no lower bound of a 95 % interval exceeds 0.24. Part of the gap between α and AC1 is the prevalence effect known from κ statistics [28, 35]. On maturity, for instance, six of the seven papers received the same consensus code, so a few dissenting votes are enough to bring α close to zero. The coders were also not fully independent: the most similar pair, `gemma4:26b` and `qwen3.8-flash-next:125b-a6b-q4_K_M`, chose the same code in a share of 0.91 of all paper–facet cells, against 0.34 for the least similar pair. Where the five local models had a majority (49 cells), the Claude pass matched it in 69 % of cases.

Of the 56 consensus codes, 42 were chosen by at least four coders and eight by all six. The remaining 14 were adjudicated against the full text of the papers, and two of them changed. One paper presents a case study without measurements, so its evidence platform is "none" rather than "target hardware" [75]. Another addresses throughput rather than latency, a tie the panel had broken alphabetically [49]. The adjudication was Claude-assisted, and both changed codes are the ones the Claude pass had chosen in the first place; we return to this dependence in section 9. The evidence for every adjudicated cell is given in the supplementary material. Table 4 is therefore best read as a description that most coders share rather than as a measurement; section 5 discusses each paper in its own terms.

5. The papers

This section takes the papers in programme order. For each we summarise the contribution as published, reproduce one figure, and report what the session added: details given only in the talk, the questions put to the presenter, and the gist of the answers. Where the published paper and the session differ, we note the difference without implying that one caused the other; the final versions were prepared after the workshop, but we hold no earlier versions to compare against. Presenters and participants are identified by role only, and quotations carry the elapsed recording time and their tier (section 2).

5.1. Faster motion detection on a microcontroller camera

Vakaliuk, Andreiev et al. [80] start from the observation that wireless transmission is the most energy-hungry operation of an IoT camera node, so a node that decides locally whether anything has moved can transmit less. On an ESP32-CAM module with an OV2640 camera, they replace normalised cross-correlation between consecutive frames with a comparison of the average brightness of blocks: a 320×240 grayscale frame is divided into a 16×16 grid, and a block counts as active when its brightness changes by more than a threshold (figure 2). One comparison of two frames took up to 105 ms with correlation and no more than 10 ms with block brightness, enough to process every frame at 25 frames per second.

The talk made the energy argument explicit: "The most energy intensive operation is data transmission to a wireless environment" (0:11:26, V). It also described a demonstration not in the paper, in which detected motion triggers a photograph sent to a messaging service. The discussion turned to conditions the laboratory test with a moving hand did not cover. A participant asked what the algorithm and its threshold should do in pouring rain, and whether frames are always converted to grayscale. The presenter explained that the camera settings had been held fixed so that the two algorithms could be compared on processing time, which does not depend on lighting, and that another camera might be chosen later. The published paper keeps this scope and names the effect of block size on detection under different lighting as a question for further research.

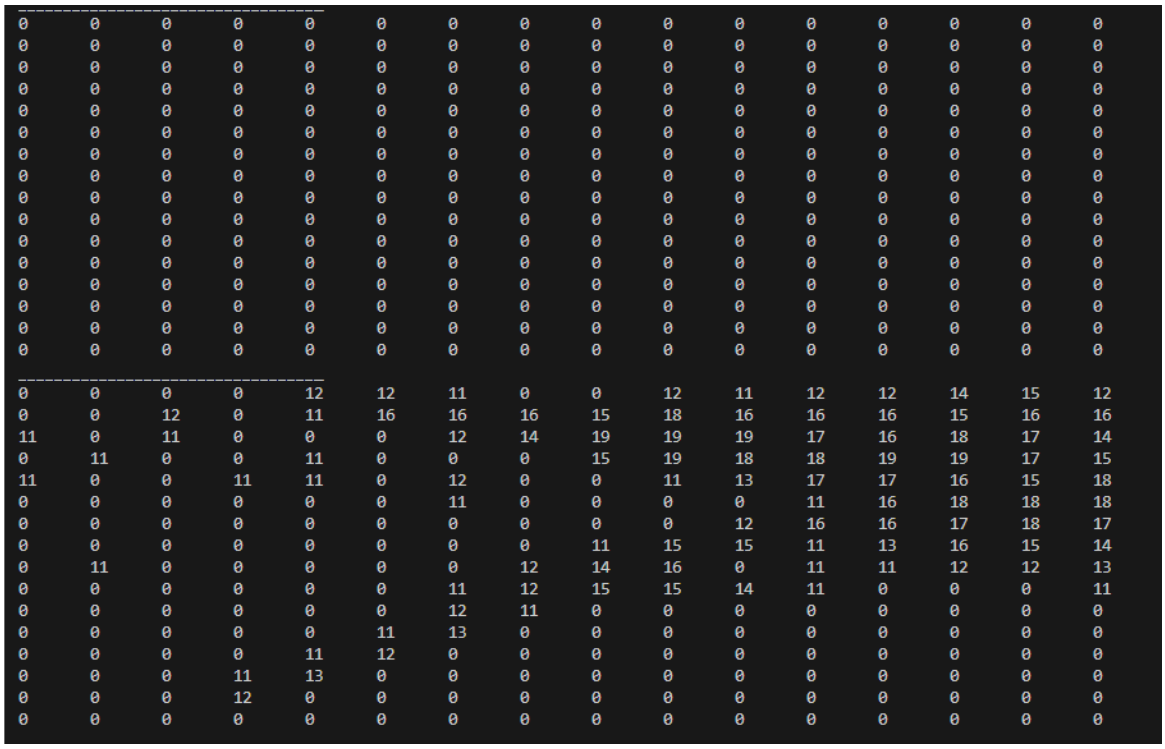


Figure 2: Differences in the average brightness ΔI_k of detection blocks: (a) without motion; (b) with motion. Reproduced from [80] under CC BY 4.0.

5.2. Event-driven perception for Feeling AI

Kargin et al. [40] apply Feeling AI, a cognitive model in which knowledge is a set of independent granules that interact through events, to the perception system of an autonomous system with distributed objects. Facts arrive asynchronously and age; the model forgets them gradually and scores their current relevance with a fuzzy certainty factor. The paper takes an automatic garage complex through every design stage, from formalising expert knowledge as granules to mapping them onto the components of an event-driven architecture and deploying those on hardware (figure 3), implemented in C++ with ZeroMQ messaging.

The presenter placed the work against generative AI: the decision core is a cognitive model, not a language model, and an everyday example of a fact that loses its relevance while the system acts motivated the ageing model. Participants asked whether the approach had been applied to smart prostheses, how granules are selected and revised, and whether the system needs a cloud service. The answers were that the laboratory platforms are wheeled robots, that granules are designed for a mission and new ones are formed at run time through the attention, emotion and drive models described in earlier publications, and that the robot performs its missions without any cloud service.

5.3. A hybrid cloud–fog architecture for industrial weighing

Alekperov et al. [6] build a three-tier system for a pharmaceutical weighing application: ESP32 microcontrollers read strain-gauge sensors, EdgeX Foundry preprocesses the data on a fog node, and InfluxDB with Grafana stores and visualises it in the cloud (figure 4). In tests with a dual-platform weighing system, a message took 155.55 ms on average from submission to storage in the cloud database, and the system processed about 440 messages per second at an offered load of 500. Local buffering kept it running through network disruptions, and the prototype hardware cost about 98 % less than commercial weighing systems.

The talk added practical detail: the fog node was an ordinary PC in the same room as the scales, HTTP was chosen deliberately to show that even a wireless link to a local fog node beats the cloud,

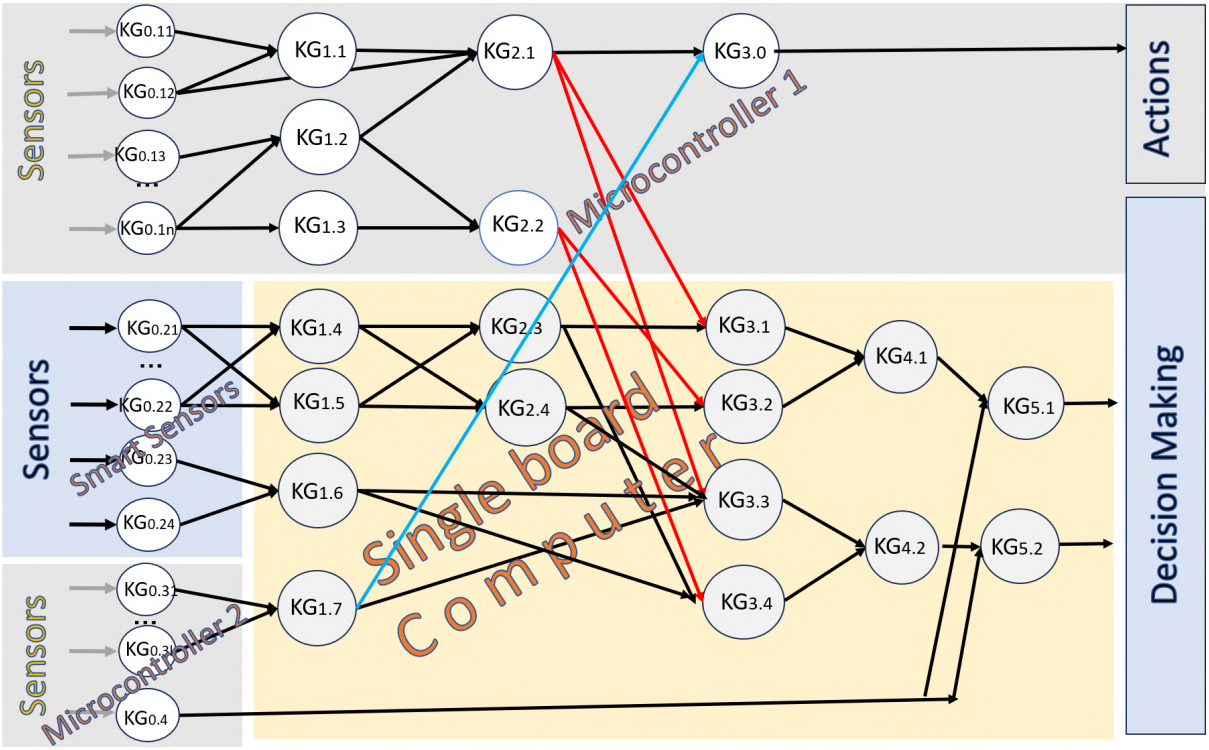


Figure 3: Mapping knowledge structure to hardware. Reproduced from [40] under CC BY 4.0.

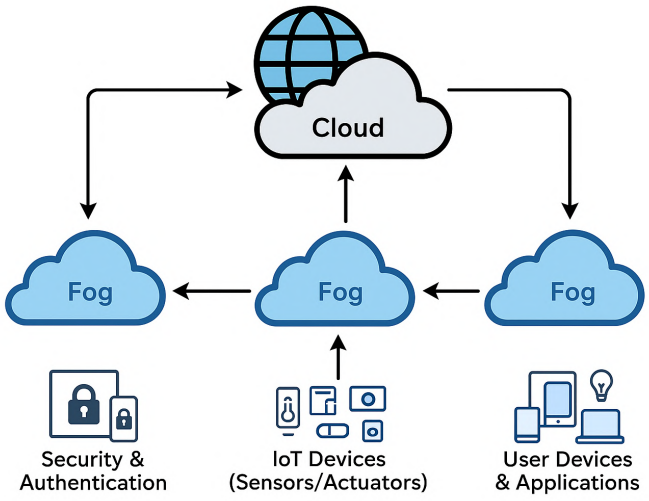


Figure 4: Hybrid cloud-fog architecture diagram. Reproduced from [6] under CC BY 4.0.

and data older than one day move to the cloud. With five asks, the discussion was the longest of the session in number of questions. One participant asked which process the reported latency and throughput referred to; the presenter described tests on two physical scales and on simulated scales, with matching results. Another participant asked three connected questions about where the statistics are computed, what share of the workload they represent and, since Grafana is a cloud service, what makes the system fog computing. The answer located the fog tier in EdgeX Foundry and the dashboard alone in Grafana. The exchange ended with the participant restating that the question concerned the boundary of fog computing itself, which remained open. The figures quoted from the slides in the discussion differ slightly from those in the published paper, which also compares the cost with commercial weighing systems rather than with the cloud IoT platforms named in the talk.

5.4. Wi-Fi parameters for camera and LiDAR streams

Karnaukh et al. [41] model a high-frequency IoT node as a pipeline from frame generation to transmission over a single Wi-Fi channel (figure 5) and study in NS-3 how frame interval, packet size and medium-access settings shape delay, jitter and losses for video and LiDAR streams. Shorter frame intervals raise delay, jitter and retransmissions; small packets add overhead and collisions, large ones are more exposed to channel errors, and among the tested sizes 900 bytes gave the best balance. Aggressive CSMA/CA settings shorten access time but destabilise the channel, while conservative modes with RTS/CTS lose less under high load.



Figure 5: Block diagram of a data transmission channel in an IoT node. Reproduced from [41] under CC BY 4.0.

The talk specified the scenario, a single-channel IEEE 802.11ac network carrying a 4K camera stream and LiDAR frames, and ended with a recommended configuration and with adaptive bitrate and the multi-user and multi-channel modes of IEEE 802.11be as future work. The four asks pressed on exactly these limits: how the results would change on several channels or under 802.11ax, what configuration an engineer should choose for a robot carrying both sensors, whether packet size could adapt to channel conditions at run time, and which kind of LiDAR was assumed. The presenter estimated the gain from multiple channels while calling it a subject for further study, said that adaptive packet size was possible but outside this study, and preferred LiDAR to cameras for detecting motion. The published paper lists the missing IEEE 802.11be capabilities among its limitations.

5.5. Few-shot species classification with frozen DINOv2 features

Lipianina-Honcharenko et al. [48] classify bird species from camera-trap images when only a few labelled examples of a rare species exist. A frozen DINOv2 vision transformer extracts features, and a k-nearest-neighbour classifier and a prototype classifier are combined by a fuzzy decision rule, so no network is trained and classification runs on a CPU (figure 6). On CUB-200-2011 and Birds525 the method reaches up to 91 % accuracy with all data and 84 % with five examples per class, above a supervised CNN ensemble and MobileNetV2, and keeps 60–65 % when half of the frame is occluded.

The talk motivated the design by the volume of unlabelled camera-trap images and by field sites that are often offline. A participant asked how much work a different taxon, such as butterflies, would take; what the caching of feature embeddings gains on an edge device; and how low-quality, night-time or blurred images would affect the embeddings, linking the last question to the one on rain in the first round. The presenter answered that a new taxon needs a new reference set rather than a new network, and that night images would need a night data set; such images had not yet been tested. The published paper reports that accuracy falls to 65–70 % on infrared night images and lists domain adaptation for them as future work.

5.6. LLM agents that configure an IoT platform

Shvaika, Shvaika and Artemchuk [75] use a multi-agent framework to turn a natural-language description of an IoT solution into a validated configuration of the ThingsBoard platform. The language model, Gemini Flash in this implementation, works only at design time, under schema-constrained reasoning; deployment is deterministic, and the generated logic consists of platform-native calculated fields and declarative alarm rules rather than imperative code. A vehicle-tracking case study takes the framework from requirement elicitation through the synthesis of entities and their topology (figure 7) to role-based dashboards operating on live telemetry.

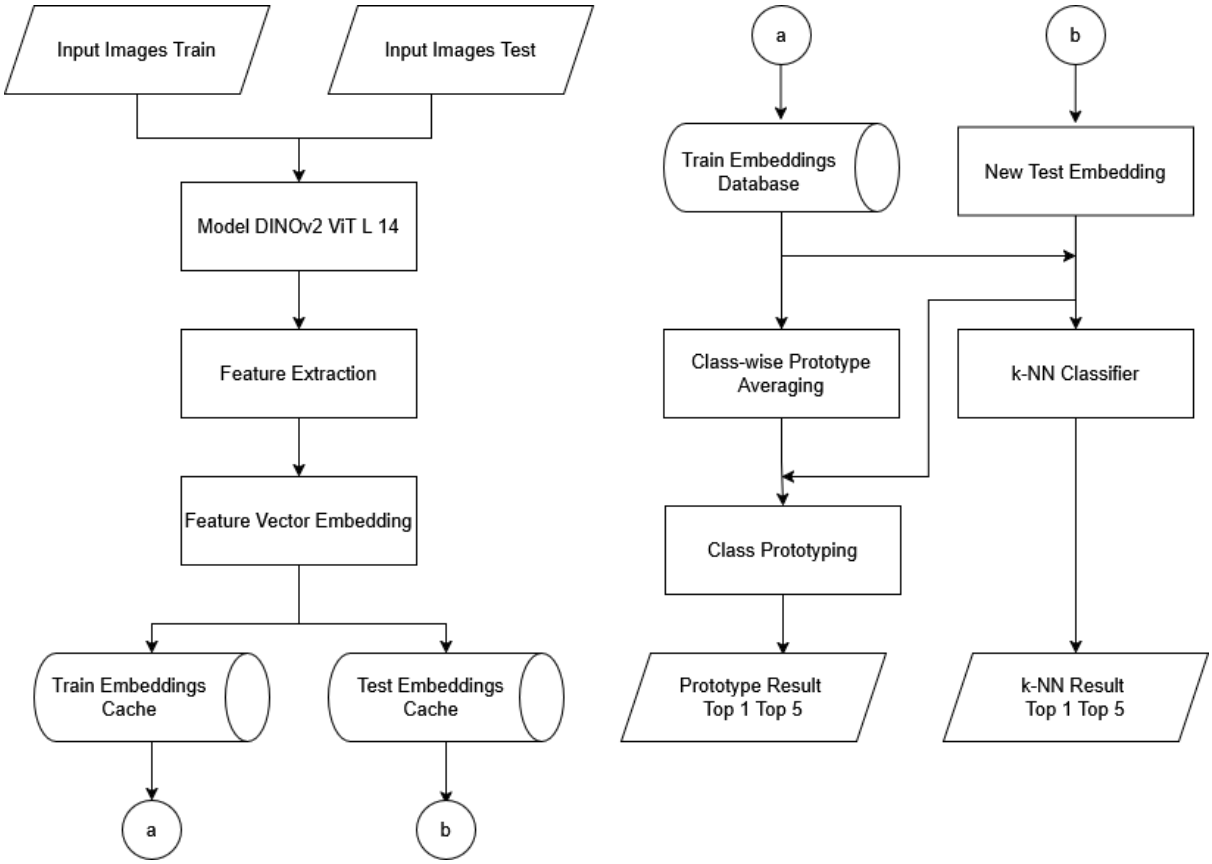


Figure 6: Architecture of the proposed hybrid classification method. The fuzzy decision rule that combines the k-NN and prototype outputs (eq.) is not drawn. Reproduced from [48] under CC BY 4.0.

This was the longest round of the session, mainly because the talk included a live demonstration and a report on work done after the paper was written: “the system now contains over 120 [...] AI [assistant] classes and over 400 prompts and schema files” (1:47:21, R). The presenter added that the framework had since been released on the vendor’s public cloud. One participant asked how the business-analyst agent was built, and the answer described its structured outputs and the assistants that consume them. Another asked whether every later change to a solution needs the language model; the answer was that the generated solution descriptor can be edited by hand and needs no model once created.

5.7. Adaptive thread pools for Python edge runtimes

Mandal and Shende [49] examine why Python AI agents on small devices lose throughput when their thread pools grow. Because the global interpreter lock serialises execution, throughput collapses beyond a certain thread count, a saturation cliff that at 2048 threads costs 40.2 % of peak throughput on one core and 35.1 % on four (figure 8). An adaptive controller driven by a blocking ratio β, which separates genuine waiting for I/O from contention for the lock, reaches 96.5 % of the best hand-tuned configuration. Multiprocessing is ruled out on small devices by an approximately eightfold memory overhead, and free-threaded Python 3.13t removes the cliff on several cores but not on one. The controller is published as an open library.

In the talk, the presenter mentioned agents running on single-board computers and phones as the motivation and explained the choice of an exponentially weighted moving average for smoothing β. A participant asked where the cliff occurs on a single core and how large the loss is, how a developer should read high and low values of β, and what the hysteresis of the controller protects against. The answers gave the range of the loss, defined β, and explained that waiting for three consecutive signals

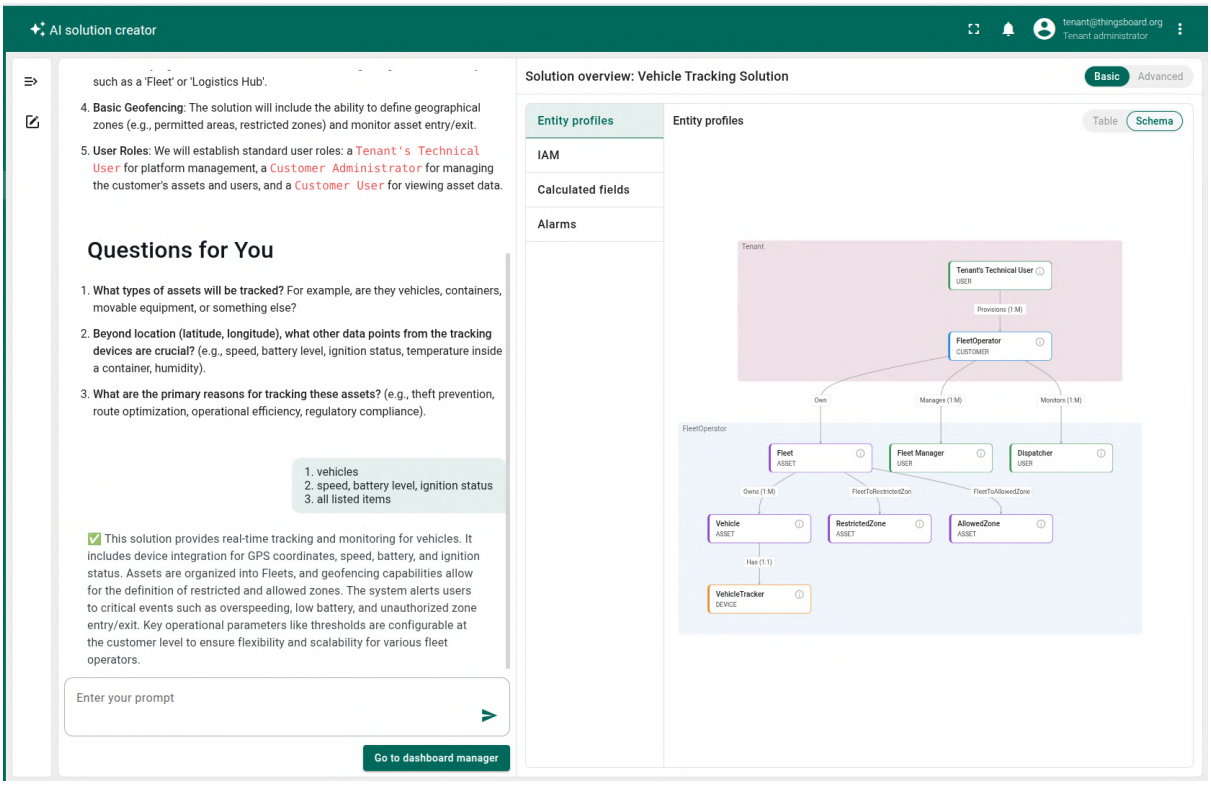


Figure 7: Automatically generated entity topology diagram derived from validated natural language requirements. Reproduced from [75] under CC BY 4.0.

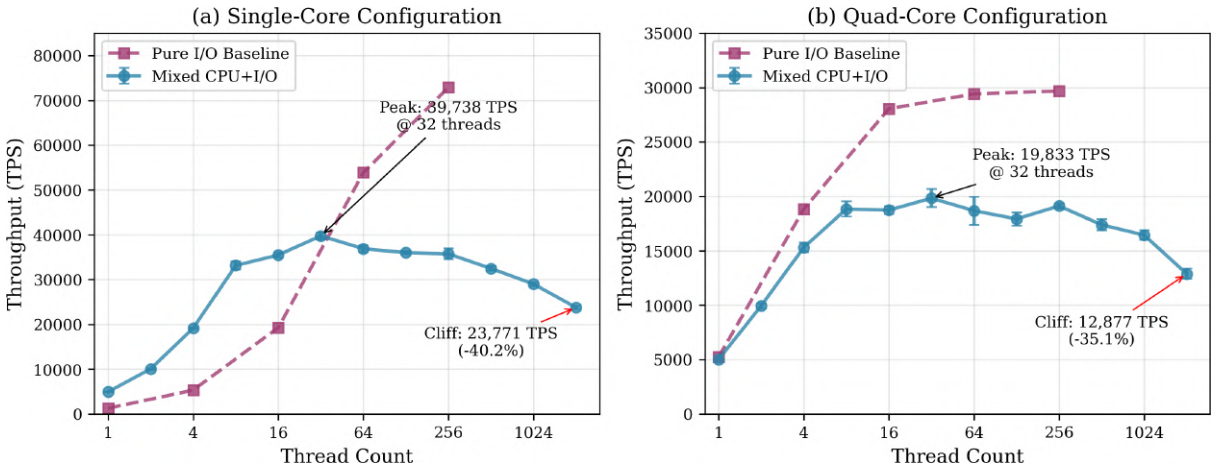


Figure 8: The saturation cliff. (a) Single-core throughput vs. threads; (b) quad-core configuration. Error bars show 95% CI ($n = 10$). Reproduced from [49] under CC BY 4.0.

stops the pool from oscillating under noisy measurements. The talk compared the controller's own small memory footprint with that of a multiprocessing pool, whereas the published paper measures the overhead of multiprocessing against a thread pool, at about eightfold; the paper also includes a workload with real inference in ONNX Runtime alongside the synthetic ones.

6. What the discussion revealed

6.1. Who asked

The 7 discussions contained 23 asks, between two and five per paper, in 43 minutes. All of them came from the floor, from five participants, three of whom were authors of other papers in the programme; the chairs moderated, called on speakers and asked for further questions, but asked nothing themselves. Most asks came from a small number of regular contributors. Seven asks opened with a word of appreciation, and four had to be repeated or rephrased.

6.2. What was asked and how it was answered

Table 5 gives the consensus codes. Most asks were information requests (19 of 23); three were challenges that presupposed a gap in the work, and one was a recommendation. Asks about method, parameters and measurement were the most common (13), followed by asks about which tier does what and whether a system is really fog or edge computing (six), and about operating conditions such as rain, night, blur and channel load (four). Deployment and cost (two), extension to new settings (three) and the dependence of a system on language models or the cloud (two) came up less often, and security, privacy and safety not at all.

Table 5
Consensus codes of the 23 asks (6 coders; majority code, ties broken by the Claude pass) with Krippendorff's α and mean pairwise Gwet's AC1. An ask can carry two topics. Codes that no ask received are omitted.

Code	Meaning	Asks
Act ($\alpha = 0.34$, AC1 $= 0.67$)		
A1	information request	19
A2	challenge	3
A3	recommendation or suggestion offered instead of a question	1
Topic (per code, see text)		
TA	operating conditions and robustness	4
TB	method, parameters and measurement	13
TC	architecture and placement	6
TD	deployment, scale and cost	2
TE	adaptation and extension	3
TF	autonomy and control of AI	2
Response ($\alpha = 0.33$, AC1 $= 0.44$)		
C1	evidence supplied	12
C2	assertion without evidence	6
C3	partial scope	2
C7	no response given	3

Answers most often supplied evidence or design detail (12 of 23). In six cases the answer was an assertion without evidence, typically an estimate or an expectation about a setting the study had not covered; in two cases it stated the limits of the study. Three asks were left without an answer to the point they raised before the turn passed on, among them follow-ups in which the questioner pressed a point after the first answer.

The coders agreed only partly, on the act (Krippendorff's $\alpha = 0.34$, 95 % CI 0.11–0.50; AC1 $= 0.67$) as on the response ($\alpha = 0.33$, CI 0.16–0.47; AC1 $= 0.44$). AC1 exceeds α on the act because almost every ask was an information request, the prevalence effect already seen in the facets. Two replies from the local models were still invalid after three attempts and are treated as missing. Topic reliability is reported only for codes with at least five positive asks, which leaves two codes. Leaving out the longest round changes α to 0.36 for the act and 0.33 for the response. The Claude pass, which extracted the asks and was not blind to the speakers, matched the majority of the local models in 81 % of the cells where they had one.

6.3. Prepared and asked questions

Before the session the chairs prepared 21 questions, three per paper, to start discussions that might otherwise not start. None was needed: every round opened with a question from the floor. Three of the prepared questions rested on a misreading, two taking the event-driven architecture of [40] for electronic design automation and one addressing the superseded title of [41]. Of the rest, 15 shared a topic with at least one ask in the same round, but only two raised the same issue as a participant did, both in the round of [6]. Both sets centred on method and measurement, but the floor raised operating conditions in four asks and placement in six, against one and two of the prepared questions.

6.4. Said and published

Because the published versions were finalised after the workshop, the discussion and the papers can be compared. In three rounds, a limitation or result raised in the discussion appears in the published paper: the effect of lighting on block-based detection is named as further research [80], the missing multi-channel modes of IEEE 802.11be are listed among the limitations [41], and night-time infrared images, untested at the time of the session, are evaluated [48]. In one round, the latency and throughput quoted from the slides differ from those in the published paper [6]. With no earlier versions to compare against, these are associations: the authors may have planned the changes before the session.

6.5. What was not discussed

Some topics were absent. No ask concerned security, privacy or safety, although several systems handle camera images, industrial data or the configuration of a multi-tenant platform. Energy, which motivated the first talk, did not come up in any discussion. Nor did reproducibility: nobody asked for code, data or the variance of a result, though only one paper links its code. These absences are a property of a single session with few active questioners, but they coincide with the gaps in the facet profile (section 4.2).

7. Positioning against the 2024–2026 literature

For each paper we collected recent peer-reviewed work on the same theme, seven or eight works per theme and 51 in all, and made two separate calls (section 2). The first concerns the contribution: whether it is *distinctive*, meaning that no direct precedent was found within the search scope, *aligned* with current practice, or *behind* it. The second concerns the evaluation: whether it is *stronger*, *comparable* or *weaker* than the evidence the retrieved works provide. Table 6 gives both calls with the works they rest on. A topic search of this size can show that a combination is rare; it cannot show that a paper is ahead of its field, and “distinctive” claims no more than the first.

Three papers have no direct precedent in the retrieved set [40, 49, 75]. Event-driven decision-making is well established in robotics, but stale observations are handled by holding the last value or by deterministic behaviour trees [12, 27], and granular computing appears mainly in resource management [33]; a graded, forgetting-based certainty for ageing events [40] is a different answer to the same problem. Its visible competitor is the small on-device language model, which a 2026 benchmark found to reach between 42.7 and 72.7 % semantic accuracy in mobile robot control [4]. Most work on language models for the IoT puts the model in the run-time loop, as an agent acting on devices [24, 61], or generates imperative code [73]; keeping it at design time and emitting declarative platform configuration [75] has its closest relative in schema-constrained generation of digital-twin models [90]. Peer-reviewed work on Python concurrency at the edge is thin: the clearest measurements of free-threaded Python are in a preprint [63], and adaptive control of thread counts has been studied outside Python [77], which leaves the adaptive thread pool of [49] without a direct comparator.

Table 6
Positioning of the seven papers against 2024–2026 work on the same theme. Contribution: *distinctive* (no direct precedent found within the search scope), *aligned* (matches current practice) or *behind* (current work has moved past the approach). Evaluation: *stronger*, *comparable* or *weaker* than the norms shown by the retrieved works.

Theme and paper	Contribution	Evaluation	Evidence
Motion detection on microcontroller cameras [80]	aligned	weaker	[3, 47, 56, 57]
Event-driven decision-making for autonomous systems [40]	distinctive	weaker	[4, 12, 27, 33, 52]
Hybrid cloud-fog architectures for industrial IoT analytics [6]	aligned	weaker	[7, 11, 17, 31, 53]
Wi-Fi parameter tuning for high-rate sensor streams [41]	behind	comparable	[1, 15, 18, 30, 39]
Few-shot species classification with frozen foundation-model features [48]	aligned	weaker	[8, 20, 26, 51]
LLM-based generation of IoT solutions [75]	distinctive	weaker	[24, 61, 73, 90]
Python concurrency for edge AI runtimes [49]	distinctive	comparable	[10, 63, 77]

Three papers follow current practice [6, 48, 80]. Deciding on the camera node whether a frame is worth transmitting is the logic of current energy-proportional vision nodes, but these use hardware triggers or quantised neural detectors rather than pixel-domain comparison [3, 47, 56]. Cloud–fog tiering with open middleware such as EdgeX Foundry is mainstream [7, 11, 17]. Frozen self-supervised features with light non-parametric heads have become a routine recipe for label-efficient ecological imaging [8, 26, 51]; the occlusion experiments of [48] address a robustness question these works leave open. One paper uses tools that current work has moved beyond [41]: sweeping packet size and contention settings on a single channel, while the Wi-Fi literature now studies multi-link operation in IEEE 802.11be [1, 30], learned channel access [18] and wake-time scheduling [15]. The joint treatment of frame interval, packet size and access settings for camera and LiDAR streams was not found elsewhere, so the study fills a narrow gap with older means.

No paper's evaluation is stronger than the norm of its theme; two are comparable [41, 49] and five are weaker [6, 40, 48, 75, 80]. The gap has a consistent shape. Where the workshop papers report processing time on a prototype, current camera nodes report energy for the whole cycle from sensing to transmission, 28.7 mJ in one 2026 design [56]. Where they observe that a fog buffer survives a disconnection, current IIoT work runs for 120 h in the field and reports tail latency and 99.92 % availability [31], or quantifies delivery reliability under injected failures [53]. Where they present a case study, the language-model literature reports success rates on task benchmarks, for example 76 % against 30 % for earlier agents [61]. And where they report a single accuracy, comparable few-shot studies report variance across seeds or confidence intervals [8, 26]. The two comparable evaluations are a simulation study, in a literature where simulation is the norm [15, 30], and the one paper that reports confidence intervals over repeated runs and links its code [49]. Measured against current work, the workshop's problems are well chosen and several of its answers are original; what it most often lacks is evidence of the kind its own field now expects.

8. Research agenda

The agenda in table 7 collects what the three lines of evidence point to together: what participants asked, what the facet profile shows the papers did and did not do, and what comparable recent work now reports. Of its seven items, six rest on all three sources; the last, on artefacts and variance, was never raised in the discussion but follows from the facet profile and the literature.

Two threads run through the items. The first is *conditions*. The papers were tested in the laboratory under nominal or synthetically varied conditions, and the questions that came up in round after round asked what happens outside it: in rain, at night, with blur, on crowded channels or when a

Table 7
Research agenda. Each item is supported by asks in the discussion, by facet codes of the papers and by recent work; the columns give the number of asks and facet codes and the works cited.

	Item	Asks	Facets	Recent work
1	Measure what on-node filtering saves [80]	2	2	[56, 57]
2	Leave the laboratory [41, 48, 80]	3	3	[26, 31]
3	Make the fog boundary explicit [6, 40]	5	2	[11, 31, 53]
4	From parameter sweeps to adaptive wireless configuration [41]	2	1	[18, 30, 39]
5	Benchmark design-time language models [75]	1	2	[61, 73, 90]
6	Validate runtimes on physical devices [48, 49]	1	2	[20, 63]
7	Release artefacts and report variance [6, 40, 41, 48, 75, 80]	–	3	[8, 26]

link fails. Recent work answers such questions with field runs, tail statistics and failure injection, and edge systems, which operate at the periphery of the network, need that kind of evidence. The second thread is *placement*. Participants repeatedly asked where the work of a system actually runs, whether the cloud is needed, whether a fog tier deserves the name, and whether a language model remains in the loop after deployment. Making placement explicit, and measuring its consequences, would answer these questions before they are asked, and would make the contributions of the workshop easier to compare with the rest of the field.

9. Limitations and threats to validity

1. *Our own roles.* We chaired doors 2026 and edit this journal, so we are analysing a session we ran for a journal we serve. T. A. Vakaliuk is also a co-author of one of the papers [80]; she took no part in coding, adjudicating or positioning that paper. We have tried to limit the influence of these roles by fixing every codebook and decision rule before applying it, by letting independently trained models do the coding, by reporting the review record only in aggregate, and by describing every paper in the same terms. The positioning verdicts in section 7 remain our judgement, and a reader who disagrees with one can check it against the works on which it rests.

2. *The recording.* Both transcripts come from automatic speech recognition, without speaker labels, of a session in which most speakers used English as a second language. Recognition errors were repaired only where both readings and the published papers agreed, and numbers were never repaired. Speakers were identified from the chairs' invitations, in some cases only on the second reading, and every ask could be attributed to a role. A recording also misses everything that happened off the record, such as conversations in the chat or after the session.

3. *The coders.* The coders are language models, not trained human annotators, and their agreement was moderate. Two of them agreed far more often than any other pair, which suggests that they share training data or design, so the panel is less independent than its size implies. The Claude pass extracted the asks and broke ties among the facet codes, and both codes changed by adjudication were the ones it had chosen in the first place. We therefore treat the coding as structured description: the counts in sections 4 and 6 are reported with the support behind them, and no conclusion rests on a code that most coders did not share.

4. *Small numbers.* Seven papers and 23 asks are too few for inferential statistics. The reliability coefficients have wide intervals, many codes occur once or twice, and the 23 asks came from only five participants, so the discussion reflects their interests as much as the community's. We report counts rather than rates where the denominators are small, and we read patterns as descriptions of this workshop, not of edge computing research in general.

5. *Positioning.* The literature search covered seven to eight recent works per theme. A search of this size can miss relevant work, so a "distinctive" verdict means only that no direct precedent was found within it. The published papers were finalised after the workshop, and we hold no earlier versions, so differences between what was said and what was published are reported as associations, not as effects of the discussion.

10. Conclusions

RQ1: the papers. The seven papers of doors 2026 share no single method but do share a place and a kind of evidence. Most of them put computation on the end device, latency is the constraint they address most often, and all of them were evaluated in the laboratory: by comparing alternatives on the authors' own data, by demonstration or by simulation, under nominal or synthetically varied conditions, with code released for one. Against work published elsewhere in 2024–2026, three papers address a combination for which we found no direct precedent, three follow current practice and one uses tools the field has moved beyond, while the evaluation of five is weaker than the norm of its theme. The evaluations, more than the ideas, are where the papers fall short of current practice.

RQ2: the discussion. The discussion took a third of the time within the rounds and consisted of 23 asks, all from the floor; the chairs asked none of the 21 questions they had prepared, three of which rested on misreadings of the papers. The questions concentrated on how results were obtained, on conditions outside the laboratory and on where computation runs, and they repeatedly anticipated limitations or extensions that appear in the published papers. Because the final versions were prepared after the workshop and we hold no earlier ones, we read this as association, not as an effect of the discussion. Most asks were requests for information rather than challenges, and most answers supplied evidence or design detail; where they did not, the question often concerned a setting the study had not covered. Security and privacy, energy and reproducibility were never raised.

RQ3: the series and the agenda. doors 2026 accepted 7 of 16 submissions, the lowest acceptance rate among the editions for which submission counts are known, in its first year in the ACM International Conference Proceedings Series. The agenda that follows from the papers, the discussion and the literature has seven items and two threads: testing under the conditions of real deployments, and making explicit, and measuring, where each part of an edge system runs. Both are questions the participants asked, and future papers can answer them in their evaluations.

11. JEC Volume 5 Issue 2 (2026): an overview

This issue of the *Journal of Edge Computing* contains six papers on pages 189–305. Besides the present review, it brings work from India, Pakistan, Bangladesh and Spain, China and Nigeria. We have ordered the papers by where their computation runs: from a server at the network edge, through edge processors and wearables, down to a microcontroller, and finally out into a deployed monitoring system.

The issue opens with this review of doors 2026 [69]. It combines the record of the sixth Edge Computing Workshop with a coded analysis of its papers and discussion, and positions the workshop against current work.

Das et al. [23] run a complete vision workflow on a 5G multi-access edge computing server to detect and classify the movements of several Kuchipudi dancers at once. A YOLO11s detector places bounding boxes. The cropped regions are downscaled, low-pass filtered in the frequency domain and contrast-enhanced, and StyleGAN2-ADA balances the seven movement classes with synthetic images. ResNet50 classifies best, with 89 % validation accuracy and a macro F1-score of 84 %. On

the CPU-only server, inference takes 44.68 ms per frame, yet the end-to-end throughput is 2 frames per second, so the stages around the model, rather than the model itself, set the pace of the pipeline.

Rohan et al. [62] detect wildfire in aerial RGB images from the FLAME 3 dataset. Their model runs a vision transformer and MobileNetV2 in parallel branches, fusing global context with local spatial detail, and Grad-CAM, Grad-CAM++ and LIME show which image regions drive each prediction. Against nine convolutional and five transformer baselines, the hybrid reaches a test accuracy of 98.44 % and a recall of 1.000 for the fire class. It classifies an image in 10–15 ms on a T4 GPU in half precision and in 50–70 ms on an x86 CPU with INT8 quantisation, which brings it within reach of edge processors carried on or near the aircraft. The authors state plainly that their results hold for the evaluated dataset and protocol.

Zuo et al. [94] bring physical reasoning into fall detection on wearable devices. PhyCL-Net adapts its temporal receptive fields to local signal dynamics with a physics-guided dynamic kernel and emphasises transient segments with jerk-aware attention; a contrastive regulariser is used during training only and removed for deployment. In leave-one-subject-out validation on SisFall, the network reaches 98.21 ± 0.10 % accuracy across five seeds with about one million parameters, and its median single-thread CPU latency is 31.6 % lower than that of a matched spectral baseline. The paper also reports where it loses: at a false-positive rate of 1 %, its true-positive rate is 2.73 percentage points lower than the baseline's.

Shah et al. [72] move gesture recognition from envelope electromyography onto a Raspberry Pi Pico. A joint search over accuracy and complexity selects the window, the overlap and four time-domain features for a small neural network, which reaches 98.35 % test accuracy and 97.22 ± 1.33 % when every participant is held out in turn. The deployed model occupies 9.1 % of flash memory and needs 23 ms per inference. In a real-time test, with a PC relaying recorded signals to the microcontroller, per-gesture accuracy over 500 consecutive predictions ranged from 95.8 to 99.6 %. The authors are careful to frame their ten-participant study as a feasibility result for prosthetic and rehabilitation interfaces.

Adesoba and Falade [2] close the issue with a system built for regions where waste collection runs on fixed schedules and bins overflow. An ESP32 with an ultrasonic sensor reports fill levels over GSM to an Android application and sends SMS alerts when a bin passes 80 % full. Across 150 sensing trials with three prototype bins, the fill estimate had a mean absolute error of 1.35 percentage points, and the application responded in 1.43 ± 0.31 s. Because acknowledgements and delivery receipts were not retained, the authors claim no transmission or delivery rates, a restraint that many system papers would do well to copy.

Read together, several of these papers supply the kind of evidence that section 7 found most often missing from the workshop: subject-independent validation [72, 94], variance across seeds [94], comparison with many baselines [62], end-to-end rather than model-only timing [23], and explicit statements of what the data cannot support [2, 62].

Author contributions

Serhiy O. Semerikov: conceptualisation, methodology, software, formal analysis, data curation, writing – original draft. Tetiana A. Vakaliuk: conceptualisation, investigation, resources, validation, writing – review and editing. Both authors chaired doors 2026 and approved the final version.

Funding

This research received no external funding.

Data availability statement

The analysis scripts, both codebooks, the coding of every rater, the adjudication and positioning

records, the session segmentation and the extracted asks with their timestamps are openly available on Zenodo at https://doi.org/10.5281/zenodo.24422591. The review record, the recording and the full transcripts contain personal data and are not shared; the review record is reported only in aggregate.

Conflicts of interest

The authors chaired doors 2026 and are editors of the *Journal of Edge Computing*. T. A. Vakaliuk is a co-author of [80] and took no part in coding or positioning it.

Acknowledgments

We thank the authors and presenters of doors 2026, the participants whose questions shaped the discussion, and the members of the programme committee for their reviews.

Declaration on Generative AI

The authors used Claude (Anthropic) to write and run the analysis scripts, to check the transcripts and segment the recording, to extract the asks from the discussion and code them as one of the raters, to adjudicate low-agreement facet codes, to search the recent literature under the rules of section 2.5, and to draft the text. The local open-weight models `gemma4:26b`, `granite4.1:8b`, `llama3.3:70b`, `nemotron3:33b`, `qwen3.8-flash-next:125b-a6b-q4_K_M` were used only as coders, through a server under the authors' control. Speech recognition used Whisper. The authors take full responsibility for the content.

References

[1] Abdalhafid, A.A., Subramaniam, S.K., Zukarnain, Z.A. and Ayob, F.H., 2024. Multi-Link Operation in IEEE802.11be Extremely High Throughput: A Survey. *IEEE Access*, 12, pp.46891–46906. Available from: https://doi.org/10.1109/access.2024.3378997.

[2] Adesoba, O.C. and Falade, S.O., 2026. Real-time smart waste monitoring using IoT technologies. *Journal of Edge Computing*, 5(2), pp.290–305. Available from: https://doi.org/10.55056/jec.1272.

[3] Adilzhanova, S., Amirkhanova, G., Kunelbayev, M. and Rakhysh, A., 2026. A Hardware Platform for Smart Video Monitoring Based on ESP32-CAM and Mojo FPGA (Spartan-6) with Event Activation Triggered by a PIR Sensor. *Engineering, Technology & Applied Science Research*, 16(3), 6, pp.36089–36096. Available from: https://doi.org/10.48084/etasr.18359.

[4] Aggul, B., Erdemir, G. and Arik, K., 2026. Benchmarking Quantized Edge LLMs for Natural Language Mobile Robot Control. *Applied Sciences*, 16(14), 7, 7267. Available from: https://doi.org/10.3390/app16147267.

[5] Alabdulwahab, S. and Moon, B., 2020. Feature Selection Methods Simultaneously Improve the Detection Accuracy and Model Building Time of Machine Learning Classifiers. *Symmetry*, 12(9), 8, 1424. Available from: https://doi.org/10.3390/sym12091424.

[6] Alekperov, V.E., Henke, K., Timenko, A.V., Kulykovska, N.A. and Diachuk, T.S., 2026. Development of hybrid cloud-fog architecture for real-time IoT analytics. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–10. Available from: https://doi.org/10.1145/3858004.3858007.

[7] Ali, T.E., Ali, F.I., Kovacs, B. and Zoltán, A.D., 2025. Comparison of IoTStack and EdgeX Foundry Performance For Asset Tracking. *Procedia Computer Science*, 258, pp.3369–3380. Available from: https://doi.org/10.1016/j.procs.2025.04.594.

[8] Azamat, U., Abdufattokhov, S., Oh, R., Ochilov, S., Kuttybayev, A., Tukhtayeva, M., Bizhanov, K., Kozhantov, A., Norinov, S., Galymzhan, S. and Kainazarova, A., 2026. GeoProtoNet: texture-aware few-shot lithology classification from rock imagery using frozen DINOv2 features. *Journal of Engineering and Applied Science*, 73(1), 9. Available from: https://doi.org/10.1186/s44147-026-01239-5.

[9] Barmak, O., Krak, I., Yakovlev, S., Manziuk, E., Radiuk, P. and Kuznetsov, V., 2024. Toward explainable deep learning in healthcare through transition matrix and user-friendly features. *Frontiers in Artificial Intelligence*, 7, 11, 1482141. Available from: https://doi.org/10.3389/frai.2024.1482141.

[10] Bartola, M.D., Massa, J., Dazzi, P. and Torquati, M., 2025. FastFlow-Python: Parallel Building Blocks in Python Through FastFlow Integration. *2025 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, pp.281–289. Available from: https://doi.org/10.1109/ic2e65552.2025.00046.

[11] Belcastro, L., Marozzo, F., Orsino, A., Talia, D. and Trunfio, P., 2026. Navigating the Edge-Cloud Continuum: A State-of-Practice Survey. *IEEE Access*, 14, pp.40622–40647. Available from: https://doi.org/10.1109/access.2026.3673012.

[12] Bonci, A., Brunella, F., Colletta, M., Biase, A.D., Dragoni, A.F. and Libofsha, A., 2026. ROS 2-Based Architecture for Autonomous Driving Systems: Design and Implementation. *Sensors*, 26(2), 1, 463. Available from: https://doi.org/10.3390/s26020463.

[13] Bonomi, F., Milito, R., Zhu, J. and Addepalli, S., 2012. Fog computing and its role in the internet of things. *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, SIGCOMM '12. ACM, pp.13–16. Available from: https://doi.org/10.1145/2342509.2342513.

[14] Braun, V. and Clarke, V., 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 1, pp.77–101. Available from: https://doi.org/10.1191/1478088706qp063oa.

[15] Busacca, F., Puligheddu, C., Raviglione, F., Rusca, R., Casetti, C., Chiasserini, C.F. and Palazzo, S., 2026. Target Wake Time Scheduling for Time-Sensitive and Energy-Efficient Wi-Fi Networks. *IEEE Transactions on Mobile Computing*, 25(3), 3, pp.3469–3487. Available from: https://doi.org/10.1109/tmc.2025.3617324.

[16] Cariow, A. and Cariowa, G., 2025. Reduced-Complexity Algorithms for Tessarine Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 36(4), 4, pp.7751–7756. Available from: https://doi.org/10.1109/tnnls.2024.3385120.

[17] Carrascal, D., Díaz-Fuentes, J., Manso, N., Lopez-Pajares, D., Rojas, E., Savi, M. and Arco, J.M., 2025. Softwarized Edge Intelligence for Advanced IIoT Ecosystems: A Data-Driven Architecture Across the Cloud/Edge Continuum. *Applied Sciences*, 15(19), 10, 10829. Available from: https://doi.org/10.3390/app151910829.

[18] Casasnovas, M., Wilhelmi, F., Combes, R., Wojnar, M., Kosek-Szott, K., Szott, S., Jonsson, A., Esteve-Elfau, L. and Bellalta, B., 2026. Learning-Based Channel Access in Wi-Fi: A Multi-Armed Bandit Approach. *IEEE Access*, 14, pp.125904–125922. Available from: https://doi.org/10.1109/access.2026.3721756.

[19] Cernohorsky, J., Diblik, M. and Richter, A., 2022. Application of motion control in rehabilitation devices. *2022 23rd International Carpathian Control Conference (ICCC)*. IEEE, pp.354–359. Available from: https://doi.org/10.1109/iccc54292.2022.9805939.

[20] Chen, Y., Zhao, D., Koh, Y.S. and Xu, T., 2026. Animal Re-Identification on Microcontrollers. *2026 IEEE International Conference on Smart Internet of Things (SmartIoT)*. IEEE, pp.249–257. Available from: https://doi.org/10.1109/smartiot70864.2026.00042.

[21] Cicchetti, D.V. and Feinstein, A.R., 1990. High agreement but low kappa: II. Resolving the paradoxes. *Journal of Clinical Epidemiology*, 43(6), 1, pp.551–558. Available from: https://doi.org/10.1016/0895-4356(90)90159-m.

[22] Dankan Gowda, V., Ingole, B.S., Agarwal, S., Prafulla P, S., Srinivas, D. and Kumari, G.S., 2024. Optimizing IoT-Based Healthcare Systems with Scalable AI and Machine Learning Using Cloud Platforms. *2024 First International Conference on Innovations in Communications, Electrical and Computer Engineering (ICICEC)*. IEEE, pp.1–7. Available from: https://doi.org/10.1109/

icicec62498.2024.10808592.

[23] Das, S., Bandaru, S., Barman, S., Jonnalagadda, A., Negi, A. and Suryadevara, N.K., 2026. Framework for real-time bounding box detection of Indian dance movements for multiple persons on a 5G multi-access edge computing server. *Journal of Edge Computing*, 5(2), 11, pp.215–232. Available from: https://doi.org/10.55056/jec.1031.

[24] Du, Y., Yang, Q., Wang, L., Lin, J., Cui, H. and Liew, S.C., 2026. LLMind 2.0: Distributed IoT Automation With Natural-Language M2M Communication and Lightweight LLM Agents. *IEEE Internet of Things Journal*, 13(23), 12, pp.61448–61462. Available from: https://doi.org/10.1109/jiot.2026.3740505.

[25] Efron, B., 1979. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1), 1, pp.1–26. Available from: https://doi.org/10.1214/aos/1176344552.

[26] Elkholy, H., Pizarro, O. and Ludvigsen, M., 2026. A comparative study of vision foundation models for few-shot classification in benthic imagery. *Frontiers in Marine Science*, 13, 11, 1925773. Available from: https://doi.org/10.3389/fmars.2026.1925773.

[27] Ernst, J., Risch, D.L., Lehner, P., Stulp, F. and Dömel, A., 2026. Decoupling what, how, and when for observing decision-making context in autonomous robots. *Frontiers in Robotics and AI*, 13, 7, 1847701. Available from: https://doi.org/10.3389/frobt.2026.1847701.

[28] Feinstein, A.R. and Cicchetti, D.V., 1990. High agreement but low Kappa: I. the problems of two paradoxes. *Journal of Clinical Epidemiology*, 43(6), 1, pp.543–549. Available from: https://doi.org/10.1016/0895-4356(90)90158-l.

[29] Ganbayev, A. and Seyidzade, E., 2023. Enhancing Customs Fraud Detection: A Comparative Study of Methods for Performance Measurement and Feature Improvement. *2023 IEEE 17th International Conference on Application of Information and Communication Technologies (AICT)*. IEEE, pp.1–5. Available from: https://doi.org/10.1109/aict59525.2023.10313153.

[30] Gao, Y., Shen, M., Zou, Y., Yin, H. and Roy, S., 2025. Latency Optimal Traffic-to-Link Allocation for MLO/SLO Coexistence in Wi-Fi 7. *IEEE Journal on Selected Areas in Communications*, 43(11), 11, pp.3633–3649. Available from: https://doi.org/10.1109/jsac.2025.3584497.

[31] García Cuadra, A., 2026. COMPASS: A multi-objective computational offloading framework for edge-cloud continuum in real-time industrial telemetry. *Computer Communications*, 257, 8, 108627. Available from: https://doi.org/10.1016/j.comcom.2026.108627.

[32] Gilardi, F., Alizadeh, M. and Kubli, M., 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30), 7. Available from: https://doi.org/10.1073/pnas.2305016120.

[33] Govea, J., Gutierrez, R. and Villegas-Ch, W., 2025. Use of granular computing for resource optimization in IoT networks. *Frontiers in Communications and Networks*, 6, 4, 1575120. Available from: https://doi.org/10.3389/frcmn.2025.1575120.

[34] Graesser, A.C. and Person, N.K., 1994. Question Asking During Tutoring. *American Educational Research Journal*, 31(1), 3, pp.104–137. Available from: https://doi.org/10.3102/00028312031001104.

[35] Gwet, K.L., 2008. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1), 5, pp.29–48. Available from: https://doi.org/10.1348/000711006x126600.

[36] Hayes, A.F. and Krippendorff, K., 2007. Answering the Call for a Standard Reliability Measure for Coding Data. *Communication Methods and Measures*, 1(1), 4, pp.77–89. Available from: https://doi.org/10.1080/19312450709336664.

[37] Hirsch, M., Mateos, C., Zunino, A., Majchrzak, T.A., Grønli, T.M. and Kaindl, H., 2021. A Task Execution Scheme for Dew Computing with State-of-the-Art Smartphones. *Electronics*, 10(16), 8, 2006. Available from: https://doi.org/10.3390/electronics10162006.

[38] Hua, X., Nikolov, M., Badugu, N. and Wang, L., 2019. Argument Mining for Understanding Peer Reviews. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, pp.2131–2137. Available from: https://doi.org/10.

18653/v1/n19-1219.

[39] Jose, P., Sukumaran, S., Suer, M.T. and Wolf, L., 2026. Can Wi-Fi 7 Meet the Latency and Reliability Demands of Time-Sensitive IIoT Networks? *2026 22nd International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. IEEE, pp.481–488. Available from: https://doi.org/10.1109/dcoss-iot69657.2026.00082.

[40] Kargin, A., Petrenko, T., Piatykop, O. and Pronina, O., 2026. AI-supported autonomous systems: EDA technology for designing Feeling AI. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–8. Available from: https://doi.org/10.1145/3858004.3858006.

[41] Karnaukh, D., Tiahunova, M., Kyrychek, H. and Kudermetov, R., 2026. Research of wireless data transmission protocol parameters for high-frequency IoT-systems based on video cameras and lidars. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–7. Available from: https://doi.org/10.1145/3858004.3858008.

[42] Kennard, N., O'Gorman, T., Das, R., Sharma, A., Bagchi, C., Clinton, M., Yelugam, P.K., Zamani, H. and McCallum, A., 2022. DISAPERE: A Dataset for Discourse Structure in Peer Review Discussions. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pp.1234–1249. Available from: https://doi.org/10.18653/v1/2022.naacl-main.89.

[43] Khatua, S., Mukherjee, A. and De, D., 2024. FedGen: Federated learning-based green edge computing for optimal route selection using genetic algorithm in Internet of Vehicular Things. *Vehicular Communications*, 49, 10, 100812. Available from: https://doi.org/10.1016/j.vehcom.2024.100812.

[44] Konzett, C., 2012. *Any Questions?: Identity Construction in Academic Conference Discussions*. DE GRUYTER. Available from: https://doi.org/10.1515/9781614510246.

[45] Krippendorff, K., 2004. Reliability in Content Analysis.: Some Common Misconceptions and Recommendations. *Human Communication Research*, 30(3), 7, pp.411–433. Available from: https://doi.org/10.1111/j.1468-2958.2004.tb00738.x.

[46] Kryzhanivskyy, V., M'Saoubi, R., Bhallamudi, M. and Cekal, M., 2022. Machine Learning based Approach for the Prediction of Surface Integrity in Machining. *Procedia CIRP*, 108, pp.537–542. Available from: https://doi.org/10.1016/j.procir.2022.03.084.

[47] Lee, C.Y., Yi, P.L., Fite, M., Rao, T., Achour, S. and Kapetanovic, Z., 2025. HyperCam: Low-Power Onboard Computer Vision for IoT Cameras. *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*, ACM MOBICOM '25. ACM, pp.529–542. Available from: https://doi.org/10.1145/3680207.3723490.

[48] Lipianina-Honcharenko, K., Stepanyk, O., Zahorodnia, D. and Bykovyy, P., 2026. Resource-efficient few-shot wildlife classification using DINOv2 and weighted fuzzy ensemble. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–10. Available from: https://doi.org/10.1145/3858004.3858009.

[49] Mandal, M. and Shende, S.S., 2026. Mitigating GIL bottlenecks in edge AI systems. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–13. Available from: https://doi.org/10.1145/3858004.3858011.

[50] Mao, Y., You, C., Zhang, J., Huang, K. and Letaief, K.B., 2017. A Survey on Mobile Edge Computing: The Communication Perspective. *IEEE Communications Surveys & Tutorials*, 19(4), pp.2322–2358. Available from: https://doi.org/10.1109/comst.2017.2745201.

[51] Markoff, H., Bengtson, S.H. and Ørsted, M., 2026. Vision Transformers for Zero-Shot Clustering of Animal Images: A Comparative Benchmarking Study. *IET Computer Vision*, 20(1), 1. Available from: https://doi.org/10.1049/cvi2.70091.

[52] Mateo-Casali, M.A., Yachouti, D.G.E., Boza, A. and Fraile, F., 2026. Analysis of Deployment Patterns for Edge Hosted Multimodal AI in Cloud–Edge Robotic Systems Utilizing ROS 2. *IEEE Access*, 14, pp.124871–124890. Available from: https://doi.org/10.1109/access.2026.3722822.

[53] Matias, J.A. and Macêdo, R.J.d.A., 2026. Resilient IIoT Communication via Edge-Based Store-and-Forward and SLO Monitoring. *Anais do XXVII Workshop de Testes e Tolerância a Falhas (WTF*

2026), WTF 2026. Sociedade Brasileira de Computação - SBC, pp.242–246. Available from: https://doi.org/10.5753/wtf.2026.24074.

[54] Matin, A., Islam, M.R., Wang, X., Huo, H. and Xu, G., 2023. AIoT for sustainable manufacturing: Overview, challenges, and opportunities. *Internet of Things*, 24, 12, 100901. Available from: https://doi.org/10.1016/j.iot.2023.100901.

[55] Milano, F., Coward, J., Jonsson, B.A., Lamm, I.L., Peterson, M., Sprawls, P., Strand, S.E. and Tabakov, S., 2009. EMITEL e-Encyclopaedia Links with EMERALD e-Learning Modules. *World Congress on Medical Physics and Biomedical Engineering, September 7 - 12, 2009, Munich, Germany*. Springer Berlin Heidelberg, pp.288–290. Available from: https://doi.org/10.1007/978-3-642-03893-8_82.

[56] Moosmann, J., Mayer, P., Benini, L. and Magno, M., 2026. An Energy-Proportional Multimodal and Context-Aware Vision IoT Node. *IEEE Internet of Things Journal*, pp.1–1. Available from: https://doi.org/10.1109/jiot.2026.3751323.

[57] Moosmann, J., Müller, H., Zimmerman, N., Rutishauser, G., Benini, L. and Magno, M., 2024. Flexible and Fully Quantized Lightweight TinyissimoYOLO for Ultra-Low-Power Edge Systems. *IEEE Access*, 12, pp.75093–75107. Available from: https://doi.org/10.1109/access.2024.3404878.

[58] Nonik, O., Lobanchykova, N., Vakaliuk, T., Osadchyi, V. and Farrakhov, O., 2024. Approaches to Solving Proxy Performance Problems for HTTP and SOCKS5 Protocols for the Case of Multi-Port Passwordless Access. *CEUR Workshop Proceedings, Vol-3654: Proceedings of the Workshop Cybersecurity Providing in Information and Telecommunication Systems (CPITS 2024)*, CPITS 2024. CEUR-WS.org, pp.189–200. Available from: https://doi.org/10.55056/ceur-ws.org/vol-3654/paper16.pdf.

[59] Querol-Julián, M. and Fortanet-Gómez, I., 2012. Multimodal evaluation in academic discussion sessions: How do presenters act and react? *English for Specific Purposes*, 31(4), 10, pp.271–283. Available from: https://doi.org/10.1016/j.esp.2012.06.001.

[60] Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C. and Sutskever, I., 2022. Robust Speech Recognition via Large-Scale Weak Supervision. Available from: https://doi.org/10.48550/ARXIV.2212.04356.

[61] Rivkin, D., Hogan, F., Feriani, A., Konar, A., Sigal, A., Liu, X. and Dudek, G., 2025. AIoT Smart Home via Autonomous LLM Agents. *IEEE Internet of Things Journal*, 12(3), 2, pp.2458–2472. Available from: https://doi.org/10.1109/jiot.2024.3471904.

[62] Rohan, S.M., Kabir, M.S., Sintheia, T.S., Tanvir, K., Gomes, D., Rahman, M. and Jony, A.I., 2026. Wildfire detection using hybrid vision transformer and MobileNetV2 with XAI. *Journal of Edge Computing*, 5(2), 11, pp.233–250. Available from: https://doi.org/10.55056/jec.1157.

[63] Salazar, J.D.M., 2026. Unlocking Python's Cores: Hardware Usage and Energy Implications of Removing the GIL. Available from: https://doi.org/10.48550/ARXIV.2603.04782.

[64] Satyanarayanan, M., 2017. The Emergence of Edge Computing. *Computer*, 50(1), 1, pp.30–39. Available from: https://doi.org/10.1109/mc.2017.9.

[65] Semerikov, S.O., ed., 2021. *Joint Proceedings of the Workshops on Quantum Information Technologies and Edge Computing (QuaInT+doors 2021)*, *Ceur workshop proceedings*, vol. 2850. Aachen: CEUR-WS.org. Available from: https://ceur-ws.org/Vol-2850/.

[66] Semerikov, S.O., 2025. Edge of arXiv 2025: bibliometrics, themes, time trends, and networks. *Journal of Edge Computing*, 4(2), 11, pp.116–158. Available from: https://doi.org/10.55056/jec.1214.

[67] Semerikov, S.O., Nechypurenko, P.P., Vakaliuk, T.A., Mintii, I.S. and Kolhatin, A.O., 2026. Energy-efficient neuromorphic computing for ultra-low latency cognitive radio: a hardware-software co-design framework for 6 G spectrum intelligence. *Discover Artificial Intelligence*, 6(1), Mar, 363. Available from: https://doi.org/10.1007/s44163-026-01093-7.

[68] Semerikov, S.O. and Vakaliuk, T.A., 2025. Edge of excellence: JEC's Scopus achievement, doors 2025 breakthroughs, and Volume 4 Issue 1 (2025) highlights. *Journal of Edge Computing*, 4(1), 5, pp.1–16. Available from: https://doi.org/10.55056/jec.987.

[69] Semerikov, S.O. and Vakaliuk, T.A., 2026. doors 2026: a quantitative review of the 6th Edge Computing Workshop, with JEC Volume 5 Issue 2 (2026) highlights. *Journal of Edge Computing*, 5(2), 11, pp.189–214. Available from: https://doi.org/10.55056/jec.1031.

[70] Semerikov, S.O. and Vakaliuk, T.A., 2026. The bio-edge: a survey and research agenda for the Internet of Bio-Nano Things, 2026-2035. *Journal of Edge Computing*, 5(1), 5, pp.1–46. Available from: https://doi.org/10.55056/jec.1382.

[71] Semerikov, S.O., Vakaliuk, T.A., Mintii, I.S., Bondarenko, O.V. and Kanevska, O.B., 2026. From science fiction to smart reality: a technology foresight review of smartdust evolution through the lens of Hamilton's *Great North Road*. *European Journal of Futures Research*, 14(1), Mar, 4. Available from: https://doi.org/10.1186/s40309-026-00266-y.

[72] Shah, M.A., Arbab, M.A., Khan, S.A. and Ullah, A., 2026. Enabling gesture recognition on a resource-constrained device: a TinyML approach with envelope EMG data and real-time testing. *Journal of Edge Computing*, 5(2), 11, pp.268–289. Available from: https://doi.org/10.55056/jec.1092.

[73] Shen, L., Yang, Q., Huang, X., Ma, Z. and Zheng, Y., 2025. GPIoT: Tailoring Small Language Models for IoT Program Synthesis and Development. *Proceedings of the 23rd ACM Conference on Embedded Networked Sensor Systems*, SenSys '25. ACM, pp.199–212. Available from: https://doi.org/10.1145/3715014.3722064.

[74] Shi, W., Cao, J., Zhang, Q., Li, Y. and Xu, L., 2016. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal*, 3(5), 10, pp.637–646. Available from: https://doi.org/10.1109/jiot.2016.2579198.

[75] Shvaika, A., Shvaika, D. and Artemchuk, V., 2026. LLM-based multi-agent framework for IoT solution generation in ThingsBoard. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–13. Available from: https://doi.org/10.1145/3858004.3858010.

[76] Song, W., Yang, Y., Liu, T., Merlina, A., Garrett, T., Vitenberg, R., Rosa, L., Awatramani, A., Wang, Z. and Birman, K., 2022. Cascade: An Edge Computing Platform for Real-time Machine Intelligence. *Proceedings of the 2022 Workshop on Advanced tools, programming languages, and PLatforms for Implementing and Evaluating algorithms for Distributed systems*, PODC '22. ACM, pp.2–6. Available from: https://doi.org/10.1145/3524053.3542741.

[77] Su, J., Hua, H., Yin, L., Yao, Y. and Zhu, F., 2026. Hybrid-Adaptive Thread Tuning to Mitigate Simulation Execution Bottlenecks in High-Performance Reinforcement Learning Inference. *Proceedings of the Thirty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI-2026. International Joint Conferences on Artificial Intelligence Organization, pp.365–373. Available from: https://doi.org/10.24963/ijcai.2026/42.

[78] Suryadevara, N.K., 2021. Energy and latency reductions at the fog gateway using a machine learning classifier. *Sustainable Computing: Informatics and Systems*, 31, 9, 100582. Available from: https://doi.org/10.1016/j.suscom.2021.100582.

[79] Vakaliuk, T.A., 2024. Edge computing research: a bibliometric analysis and JEC Volume 3 Issue 2 (2024) highlights. *Journal of Edge Computing*, 3(2), 11, pp.110–125. Available from: https://doi.org/10.55056/jec.853.

[80] Vakaliuk, T.A., Andreiev, O.V., Dubyna, O.F., Nikitchuk, T.M. and Shevchenko, M.V., 2026. Increasing the speed of digital image processing in edge devices with limited computing resources. *Proceedings of the 6th Edge Computing Workshop*, doors 2026. ACM, pp.1–5. Available from: https://doi.org/10.1145/3858004.3858005.

[81] Vakaliuk, T.A., Chyzhmotria, O.V., Chyzhmotria, O.H., Didkivska, S.O. and Kontsedailo, V.V., 2023. The use of massive open online courses in teaching the fundamentals of programming to software engineers. *Educational Technology Quarterly*, 2023(1), 1, pp.106–120. Available from: https://doi.org/10.55056/etq.37.

[82] Vakaliuk, T.A., Nikitchuk, T.M., Andreiev, O.V., Medvediev, M.G. and Matsiievskiy, V., 2026. Integrated Hybrid Biomedical Sensor Network for Multi-layered Telemedicine Monitoring. *Sensors, Devices and Systems*. Springer Nature Switzerland, pp.264–273. Available from: https://doi.org/10.1007/978-3-032-23369-1_27.

[83] Vakaliuk, T.A. and Semerikov, S.O., 2023. Introduction to doors Workshops on Edge Computing (2021-2023). *Journal of Edge Computing*, 2(1), 5, pp.1–22. Available from: https://doi.org/10.55056/jec.618.

[84] Vakaliuk, T.A. and Semerikov, S.O., eds, 2023. *Proceedings of the 3rd Edge Computing Workshop*, *Ceur workshop proceedings*, vol. 3374. Aachen: CEUR-WS.org. Available from: https://ceur-ws.org/Vol-3374/.

[85] Vakaliuk, T.A. and Semerikov, S.O., 2024. Empowering the Edge: Research advances from doors 2024. *Journal of Edge Computing*, 3(1), 5, pp.3–27. Available from: https://doi.org/10.55056/jec.747.

[86] Vakaliuk, T.A. and Semerikov, S.O., eds, 2024. *Proceedings of the 4th Edge Computing Workshop (doors 2024)*, *Ceur workshop proceedings*, vol. 3666. Aachen: CEUR-WS.org. Available from: https://ceur-ws.org/Vol-3666/.

[87] Vakaliuk, T.A. and Semerikov, S.O., eds, 2025. *Proceedings of the 5th Edge Computing Workshop (doors 2025)*, *Ceur workshop proceedings*, vol. 3943. Aachen: CEUR-WS.org. Available from: https://ceur-ws.org/Vol-3943/.

[88] Vakaliuk, T.A. and Semerikov, S.O., 2026. Edge Computing in Practice: Contributions and Conversations at doors 2026. *Proceedings of the 6th Edge Computing Workshop*, doors 2026, workshop report in the proceedings front matter. ACM, pp.iv–x. Available from: https://doi.org/10.1145/3858004.

[89] Wulff, S., Swales, J.M. and Keller, K., 2009. "We have about seven minutes for questions": The discussion sessions from a specialized conference. *English for Specific Purposes*, 28(2), 4, pp.79–92. Available from: https://doi.org/10.1016/j.esp.2008.11.002.

[90] Xia, Y., Xiao, Z., Jazdi, N. and Weyrich, M., 2024. Generation of Asset Administration Shell With Large Language Model Agents: Toward Semantic Interoperability in Digital Twins in the Context of Industry 4.0. *IEEE Access*, 12, pp.84863–84877. Available from: https://doi.org/10.1109/access.2024.3415470.

[91] Yilmaz, O., Gorgu, L., O'Grady, M.J. and O'Hare, G.M.P., 2021. Cloud-Assisted Mobile Crowd Sensing for Route and Congestion Monitoring. *IEEE Access*, 9, pp.157984–157996. Available from: https://doi.org/10.1109/access.2021.3129932.

[92] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K. and Zhang, J., 2019. Edge Intelligence: Paving the Last Mile of Artificial Intelligence With Edge Computing. *Proceedings of the IEEE*, 107(8), 8, pp.1738–1762. Available from: https://doi.org/10.1109/jproc.2019.2918951.

[93] Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z. and Yang, D., 2024. Can Large Language Models Transform Computational Social Science? *Computational Linguistics*, 50(1), pp.237–291. Available from: https://doi.org/10.1162/coli_a_00502.

[94] Zuo, B., Xie, J., Wang, S., Liu, S. and Luo, X., 2026. Physics-guided contrastive learning for wearable fall detection: towards low-latency on-device inference. *Journal of Edge Computing*, 5(2), 11, pp.251–267. Available from: https://doi.org/10.55056/jec.1262.