

Artificial Situational Awareness Assessment of a Novel ATC Support System

I Tukarić^{1,3}, K Samardžić¹, T Radišić¹, C Vetter²

¹ Department of Aeronautics, Faculty of Transport and Traffic Sciences, University of Zagreb, Vukelićeva 4, 10000 Zagreb, Croatia

² School of Engineering, Centre for Aviation, Zürich University of Applied Sciences, Technikumstrasse 9, 8401 Winterthur, Switzerland

³ itukaric@fpz.unizg.hr

Abstract. This paper presents the application of an existing situational awareness framework to a newly developed artificial intelligence system to determine its awareness level. The system incorporates diverse automation techniques - knowledge graph, expert rules, machine learning - for gaining situational awareness and applying it in the field of air traffic control. Since the system was developed to serve as a foundation for exploring automation and artificial situational awareness, the primary result of this work is the system's overall awareness level assessment and the identification of sub-systems that may be improved for additional awareness. The framework used was chosen in the fundamental project documents and its use proved beneficial as it enabled the demonstration of how general guidelines can be interpreted for a specific system. It also informed possible routes for improvement of the process. Highest priority awareness-related improvements are those dealing with robustness, whose implementation would substantiate the current awareness assessment. The system is shown to be on the highest awareness level conditionally, considering its proof-of-concept level. The high level reached by the system is contingent on awareness concept and condition interpretations. With the appropriate assessment of the system, implementation in an operational environment is more feasible.

Keywords: artificial situational awareness; air traffic control; knowledge graph; machine learning

1. Introduction

The push for automation in air traffic management (ATM) has been spurred by the Single European Sky ATM Research Joint Undertaking (SESAR JU) – their European ATM Master Plan envisions higher levels of automation as support for the eventual digitalization of the European sky [1]. The document proposes five levels of automation, ranging from “Decision support” to “Full automation”, as general guidance for projects under the SESAR banner. The projects are accompanied by wider systemic efforts to standardize information exchange in a modern way, e.g., by employing information exchange standards (Aeronautical Information Exchange Model/AIXM, Flight Information Exchange Model/FIXM).

One of SESAR projects is the AI Situation Awareness Foundation for Advancing Automation (AISA), aimed at introducing a foundation for automation by first developing a situationally aware system [2].



In this context, situational awareness (SA) is defined as the “perception of environmental elements and events with respect to time or space, the comprehension of their meaning, and the projection of their future status.” [3]. A system matching this description could then be used to promote human-machine distributed situational awareness in, for example, en-route air traffic control (ATC) tasks.

Development of human-level SA in an artificial intelligence system is still an unattainable goal. This is due to their inability to mimic human neural functions and achieve cooperation between different sub-systems that is characteristic for the human brain (known as effective connectivity) [4]. Other paths to artificial SA must therefore be found, which proved to be an interesting problem for a plethora of projects.

Measuring and determining human SA with several approaches to it can be as perplexing as artificial SA. The construct of human SA itself makes assessing a complex argument. The lack of a broadly accepted model of SA makes it difficult to select a measure that is consistent with the construct definition. One of the methods is the individual's *subjective opinion* of their own SA. This method is the most direct and doesn't require many resources but lacks validity as an individual can assess his SA only to a certain level and cannot assess SA on the aspects he is not aware of. Thus, it is rather an opinion participants have on their own SA [5].

Another method is *implicit performance measurement* where SA is measured based on some predefined performance indicators [6]. Probe techniques ask the participant questions about the current traffic situation and the one that will evolve. This method can be performed in two ways - probe technique with freezing includes stopping the simulation and asking questions while blinding information sources (such as radar screens) and the online probe technique where questions are asked during task performance, possibly contributing to a higher workload or performance obstruction.

Process indices record the examinees' performance and analyse it during task execution. An example of the SA process indices is eye-tracking techniques with simultaneous verbal protocol analysis. This method can determine how the attention was deployed during task performance (eye-tracking) with an available written transcript of the participants' actions during task execution (verbal protocol). The transcript is used for getting perception into the cognitive aspects of participants' complex actions [6]. All presented methods cannot be used in isolation in the assessment of human SA (real or controlled environment). Although some methods are more common, the struggle of defining the exact aspects of SA that can be assessed for individual and team SA remains an obstacle [7].

The proof-of-concept system developed during the AISA project employs different automation techniques to achieve artificial SA. Its architecture is shown in Figure 1. The effectiveness of the system has already been established by comparing system outputs with ATCO SA (using all listed SA measurements), but a standalone artificial awareness analysis is also deemed necessary. Aside from

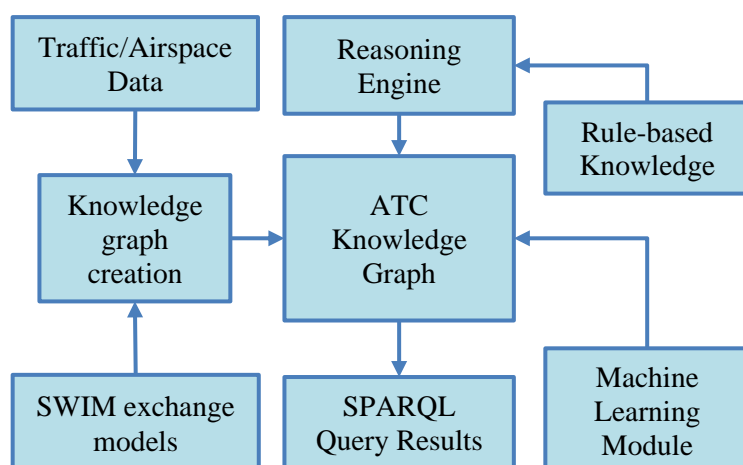


Figure 1. Simplified architecture of the AISA System

taking part in distributed SA, artificial awareness is also named as a way to improve system robustness by Jantsch and Tammemäe [8].

The same article presents a classification system for artificial intelligence in highly resource constrained systems, with the aim of studying awareness implementation methods. The question is how can we evaluate and determine the system's overall awareness level and identify sub-systems that may be a subject of improvement for additional awareness? The framework presented in [8] proposes how the AI system with its components can be classified

in such a way that general guidelines are applied and designed to suit specific systems. Also, these guidelines can identify potential obscurities during the process and possible routes for improvement. This article examines how that framework can be applied to the AISA system, which awareness level it reaches, and how framework requirements can be used to further improve the system's awareness.

2. Awareness level classification framework

As with human SA measurement, artificial SA can also be measured and assessed using several established methods and approaches. Artificial SA assessment can be solely an expert's opinion on the generated artificial awareness [9]. Mitchell presented a list of four principles that interact and that are applicable for distributed self-aware system assessment [10]. Self-aware node is a conceptual component that is locality within a global system, not especially existing as a hardware or software. This node has several conditions where the information about internal state (private self-awareness) and state of its environment (public state-awareness) is obligatory to be met. Optionally, node can have information of its role or importance, effect of future actions and historical knowledge. Even though this research does identify what are the conditions of self-aware system, those conditions are hardly applicable for the assessment of the system's situational awareness.

When discussing self-awareness, connections to human psychology are inevitable. The comparison of self-awareness and self-expression in cognitive science with artificial systems identified five levels of self-awareness: (i) ecological self (minimum requirements for an object to not be unconscious), (ii) interpersonal self (a simple awareness with limited adaptation), (iii) extended self (object is aware of the past and future actions), (iv) private self (object can process more information about its state), (v) conceptual self (object can construct a symbolic representation of itself) [11]. This definition allows scalability of system's complexity and self-awareness properties (direct or emergent).

A more objective approach involves the use of framework that, based on the defined conditions, evaluates the system's awareness level and ability of the system to assess its own performance [8]. The aforementioned article by Jantsch and Tammemäe presents a six-level classification system for artificial intelligence in highly resource constrained systems, with the aim of studying awareness implementation methods. This article examines how their framework can be applied to the AISA system by analysing the fulfilment of each awareness condition and requirement necessary to achieve the described awareness levels.

Project AISA's concept of operations proposed the use of an existing framework for AI system awareness level assessment [2]. Framework elements and their mutual dependencies are shown in Figure 2. The chosen framework consists of:

- a vocabulary,
- awareness conditions,
- awareness requirements,
- a classification system.

The vocabulary is made up of 11 terms deemed necessary to accurately define awareness related capabilities and awareness levels. Those terms include abstraction (which is the process of mapping measured data to values of a selected property from a set), semantic attribution (which is the mapping of property values to a point on a desirability scale), etc.

The awareness conditions are divided into 2 categories: i) conditions for the awareness of a property and ii) conditions for the awareness of the system itself. The authors recognize that the system may be aware of the

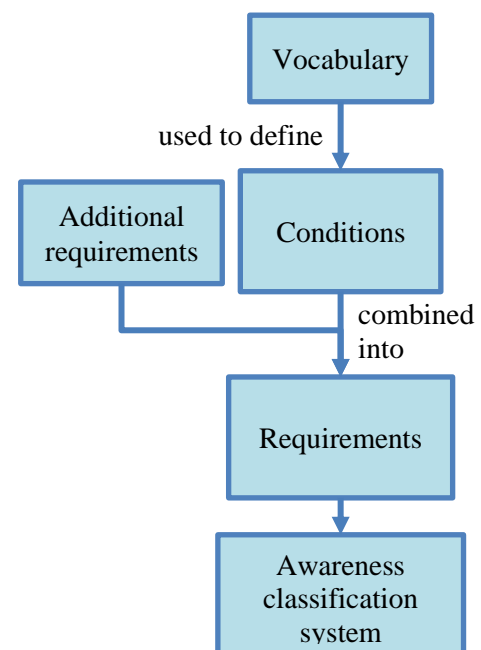


Figure 2. Framework elements and layout

environment, itself, or both, but it can never reach complete knowledge of either. This, however, is not necessary – awareness does not depend on information volume.

A combination of awareness conditions and additional clauses pertaining to specific sub-systems or system functions makes up the **requirements** for reaching each level of awareness in the **classification system**. The number of conditions for a certain level is either the same or greater than for the preceding one, but the additional constraints are always different between levels.

3. Classification of AISA system

To achieve artificial SA, the AISA system uses an ATC knowledge graph (KG), developed specifically for the project. While the KG is mainly based on concepts defined in the information exchange models (AIXM FIXM), they are expanded with project-specific concepts. The KG stores traffic situation data, organized according to a predefined structure and rules, so it can be used to generate predictions via machine learning (ML) modules and to generate SA via an expert-knowledge-based reasoning engine. Additionally, the KG is a representation of all necessary attributes, rules, relationships, axioms, etc. from the ATC field. It can therefore be used as a basis for expanding the SA by gaining new knowledge and drawing conclusions about the state of the system/environment, in parallel with the reasoning engine.

The original awareness framework article does not offer a methodology for applying it to existing AI systems. The conditions and awareness levels will therefore be used as guidelines. Preliminary analysis shows that an informed choice regarding system scope must be made since it will influence the classification results. Specifying that scope is not always a straightforward task – it depends on the system itself and the perspective of the person performing the analysis. For the purposes of classification of the AISA system, a system “border” is set between the AISA system (consisting of a KG, reasoning engine/AISA tasks, and the ML modules) and the simulator on the input side and ATCO on the output side. This is presented in Figure 3.

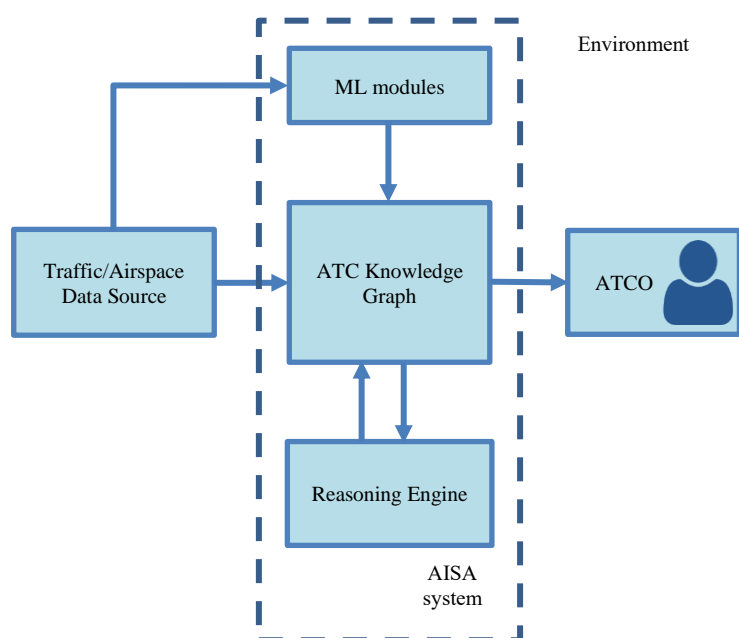


Figure 3. AISA system scope for the purpose of awareness assessment

Since the AISA system is at a proof-of-concept level and not working in real-time, connections between sub-systems and the environment are not as they would be in a fully developed ATC system. That future system could be seen as having ATCOs and even raw data sources included within the system border, with pilots outside of the environment. Those changes in the system border could then lead to a different awareness level estimation.

Conditions defined in the framework and listed earlier in this document will be analysed from the perspective of the system and exemplified. Fulfilment of these conditions will then be used to classify the AISA system on the awareness level scale.

3.1. Fulfilment of awareness conditions

C.1 Meaning condition: *Subject makes physical measurements or observations that are used to derive the values of property P by means of a meaningful semantic interpretation.*

Measurements and observations from which values of a property can be derived are gathered from the various data sources the AISA system has at its disposal. For example, the simulator delivers the values of flight information such as flight level, speed, position, and others. Although the system itself is not performing the measurements and observations, it is gaining values of properties from the environment in the process of data translation.

A meaningful semantic interpretation consists of mapping gathered measurements to values of a property and choosing the appropriate interpretation if mapping results in more than one interpretation. Since the conversion of values (from data files to Resource Description Framework (RDF) graphs) maps values to properties directly, only a single interpretation is possible. The properties (flight level, speed etc.) are, of course, meaningful in the context of ATC, so all parts of condition C.1 are fulfilled.

C.2 Robustness condition: *The semantic interpretation is robust.*

The robustness of semantic interpretation is the task of SHACL rules, which validate RDF graphs against a set of conditions. In case of faulty inputs - originating, in this case, from either human error or the program tasked with converting data to RDF format - the system returns an error and points to where the error occurred and why. Since the system uses both existing and user-defined concepts as the basis for the KG, SHACL rules may only detect some, instead of all, erroneous inputs. Therefore, robust semantic interpretation is accomplished, but *not fully guaranteed*.

C.3 Attribution condition: *There is a semantic attribution which is meaningful.*

Semantic attribution - mapping values of a property to a desirability scale - is performed by the automated ATCO monitoring tasks by comparing actual property values to those defined by the goals (e.g., cleared values). Values are then implicitly graded as desirable (i.e. equal to cleared) or not desirable. The desirability is expressed through monitoring task outputs. Not all values are mapped to a desirability scale – this is deemed to be acceptable because desirability (beyond the base test that are the SHACL rules, which are already applied to all properties) cannot be established for properties such as callsigns or statistical values of conflict ML module training data.

An example of a *system* property being checked for desirability is the system's inspection of the conflict prediction ML module operation. It is checked for desirability of input data by comparing it to the statistics of the module's training set. The deviation of the actual value from the mean value of the training set is visible in the task output.

C.4 Appropriateness condition: *The subject's reaction to its perception of P is appropriate.*

The AISA system achieves appropriate reaction to the perception of properties by:

- analyzing and storing the values of properties
- using values of properties for creating other properties and computing their values
- creating appropriate outputs for values of properties

C.5 History condition: *A history of the evolution of the property over time is maintained, in particular of the increasing or decreasing deviations over time.*

History of evolution of each property is easily accessible since each situation graph is stored in the KG, along with output graphs collecting all monitoring tasks outputs. By using SPARQL queries it is possible to access the graph of each timestamp and retrieve the value of any property. A similar effect can be accomplished, while decreasing memory necessary for storing previous situation graphs, by developing a conversion mechanism which would create graphs representing the previous property values while using less space in the KG.

Increasing or decreasing deviations can be tracked through task outputs – they are implicit in outputs such as “Aircraft is not at CLFL.” and “Aircraft is descending towards CLFL.”

C.6 Goal condition: *The subject can assess how well it meets all its goals, thus having an understanding which goals should be achieved and to which extent they are achieved.*

Goals of the AISA system are represented within the KG via cleared values. Coupled with the automated tasks, the KG can check and state which goals are achieved (e.g. “Aircraft is at cleared speed.”) or are currently being achieved (e.g. “Aircraft is climbing towards CLFL.”). Nature of cleared values in ATC means that all goals described this way in the KG must be achieved, to maintain the safety of air traffic.

C.7 Goal History condition: *The subject can assess how well the goals are achieved over time and when its performance is improving or deteriorating.*

The AISA system runs all automated ATCO monitoring tasks and can, by analysing the outputs, check the status of each goal and its changes through the scenario. The storing of task outputs ensures goal completion can be assessed over time. The system also works through the ATCO to ensure all goals are being achieved, without preference. As an example of direct performance monitoring, tasks related to the operation of the conflict detection ML module monitor both the status of each conflict (which are some of the goals of the system) and the performance of the module itself (the correctness of each prediction). In this way, the system checks both the goals and its own performance.

3.2. Awareness level assessment

The framework identified six levels of AI system awareness. The levels are shown here through their requirements, together with how the AISA system fulfils them. Since fulfilment hinges on awareness condition fulfilment, it is important to summarize the previous section and state that all conditions are fulfilled, except for condition C.2, which is partially fulfilled.

Table 1. AISA awareness level estimate

Awareness Level	Necessary requirements to reach level	AISA KG system function
Awareness Level 0	<ul style="list-style-type: none"> • System output is a mathematical function of inputs (always reacting in the same way to inputs) • System fulfils conditions C.1 to C.4 	<p>The AISA system consists of computer code which, for identical inputs, always produces the same output.</p> <p>Conditions C.1, C.3, and C.4 have been shown to be fulfilled earlier in this chapter. Condition C.2 is partially fulfilled (since it’s not guaranteed), so Awareness Level 0 requirements can be thought of as <i>partially fulfilled</i> as well.</p>
Awareness Level 1	<ul style="list-style-type: none"> • System is adaptive, meaning that it tries to minimize the difference between input and reference values by use of a PID controller or similar algorithm • System fulfils conditions C.1 to C.4 	<p>The AISA system fulfils the adaptiveness condition by having the outputs of the KG system point toward the difference between actual and goal values. That affects the actions of the ATCO, thus ensuring the minimisation of differences.</p> <p>Conditions C.1, C.3, and C.4 have already been shown to be fulfilled by the AISA system. Condition C.2 is partially fulfilled (since it’s not guaranteed) so Awareness Level 1 requirements can be thought of as <i>partially fulfilled</i> as well.</p>
Awareness Level 2	<ul style="list-style-type: none"> • System is aware of at least one (system) property and one environment property according to C.1 to C.4 + 	<p>The system is aware of both environment properties (such as A/C trajectories) and system properties (such as conflict detection module performance) in ways prescribed by requirements of this level: property values are derived from</p>

C.6

- System contains an inspection engine which periodically derives one integrated attribution of the system as a whole
- System computes its actions based on (a) monitored and attributed properties of the system and of the environment, (b) attributed expectations on the system and on the environment, and (c) sets of goals on system and environment properties

measurements, the semantic interpretations are robust (in ways defined in the conditions), the property values are checked for desirability (via automated tasks) and task outputs are linked to those values (and are therefore appropriate).

The inspection engine condition is fulfilled by the *conflict detection module* – tasks which check the desirability of module inputs (against training data statistics) and outputs (by way sanity check and basic comparison calculations) are a way for the system to analyse its own behaviour.

The AISA system computes necessary actions according to the property values defined in the KG (such as the already mentioned A/C FLs or conflict detection module performance), expectations on itself and the environment (which are defined by SHACL rules and KG completeness), and goals (which are contained in the KG). Since the expectations on the system are contingent on functioning of SHACL rules, this condition and awareness level cannot be guaranteed to be fulfilled. For this reason, Awareness Level 2 is reached *partially*.

Awareness Level 3

- System fulfils all requirements of an Awareness Level 2 system
- System fulfils the history conditions C.5 and C.7

The history conditions are fulfilled as demonstrated in section 0– each timestamp’s traffic data and task output graphs are stored in the KG and easily accessible. Combined, they form a history of each property and property value where values are direct proof of deviations. The improvement and deterioration are demonstrated only for appropriate properties – e.g. conflict detection module performance.

Fulfilment of Awareness Level 2 requirements is shown in the cell above. Since that level is reached only partially, Awareness Level 3 is also reached *partially*.

Awareness Level 4

- System fulfils all requirements of an Awareness Level 3 system
- System decision-making process involves a simulation engine which can predict the effects of actions on the environment and the system itself and, in case of an anomalous result, can search through simulations for the best action

Simulation engine requirement is completed by the machine learning module, which uses each traffic data graph as input and calculates how modifications of certain property values can lead to different traffic outcomes. A voluntary number of repetitions (with unique value modifications) can be performed, and the results parsed for the optimal action (or actions).

Since the achievement of this awareness level hinges on the achievement of Awareness Level 3, it is deemed to be *partially achieved*.

Awareness Level 5

- In addition to being self-aware, the system distinguishes between itself, environment, and peer group (treated differently because of its own set of expectations and goals)

The AISA system contains tasks dealing with environment and system properties, but also with properties formed by third parties (such as sector exit flight levels, dictated by agreements with neighbouring air navigation service providers). Those providers can be seen as a peer group with specific goals, whose existence is recognized by the KG. Thus, Awareness Level 5 is seen as *partially achieved*.

4. Conclusions and future work

According to the chosen framework, the AISA system is conditionally an Awareness Level 5 system. The work presented herein demonstrates a possible classification process for AI systems, since one was not presented in the original work. The process consisted of choosing a system boundary and then identifying which system component best suits the awareness conditions and requirements. This was complicated by the fact that the AISA system does not run in real-time and its sub-systems are not fully integrated. Table 1. describes how the AISA system achieves each awareness levels by highlighting only the best-performing sub-system and using it to represent the whole system. This approach is deemed acceptable since the use of specialized sub-systems is a common strategy in biological and technological systems.

The fact that the AISA system reaches Awareness Level 5 only conditionally depends solely on the current method of checking system inputs against constraints - SHACL rules. If the current method is expanded to include all inputs or supported by a second layer of checks, the awareness level assessment could be confirmed. A future functional system could improve or replace insufficient sub-systems that degrade the overall awareness level so that a higher degree of awareness can be assigned to it. This demonstrates how the application of this framework helps identify awareness “weaknesses” in the system, which are then prime candidates for future work and research.

The awareness assessment could thus benefit from more flexible language regarding setting system borders and fulfilment of specific awareness requirements. If the language was clearer but more restrictive, it would preclude the application of this framework to many automated systems, including the AISA system. Specifics of framework application are best left to system creators, but the authors agree that choosing the best-performing sub-system is a valid approach.

Acknowledgments

This paper is part of the AISA project, which has received funding from the SESAR Joint Undertaking under Grant Agreement No 892618 under European Union’s Horizon 2020 research and innovation program.

References

- [1] SESAR JU. European ATM Master Plan, <https://www.sesarju.eu/masterplan2020> (2020).
- [2] AISA Consortium. Concept of Operations for AI Situational Awareness (D2.1)., https://aisa-project.eu/downloads/AISA_D2.1_CONOPS.pdf (2020).
- [3] Endsley MR. Toward a Theory of Situation Awareness in Dynamic Systems. *Hum Factors* 1995; 37: 32–64.
- [4] Thiopoulos C. Can AI Systems Match Human-Level Situational Awareness? *Mouser Electronics*, <https://www.mouser.com/blog/can-ai-systems-match-human-level-situational-awareness> (2020, accessed 25 May 2022).
- [5] Endsley MR. Measurement of situation awareness in dynamic systems. *Human Factors* 1995; 37: 65–84.
- [6] Salmon PM, Stanton NA, Walker GH, et al. Measuring situation awareness in complex systems: comparison of measures study. *International Journal of Industrial Ergonomics* 2009; 39: 490–500.
- [7] Salmon P, Stanton N, Walker G, et al. Situation awareness measurement: a review of applicability for C4i environments. *Applied Ergonomics* 2006; 37: 225–238.
- [8] Jantsch A, Tammemäe K. A framework of awareness for artificial subjects. 2014.
- [9] Pullum L. *Verification and Validation of Systems in which AI is a Key Element*. Hoboken, New Jersey, United States of America: Oak Ridge National Lab., 2021.
- [10] Mitchell M. Self-Awareness and Control in Decentralized Systems.
- [11] Lewis PR, Chandra A, Parsons S, et al. A Survey of Self-Awareness and Its Application in Computing Systems. In: *2011 Fifth IEEE Conference on Self-Adaptive and Self-Organizing Systems Workshops*. Ann Arbor, MI, USA: IEEE, pp. 102–107.