World Scientific Research
Vol. 11, No. 1, 6-11, 2024
ISSN(E) 2411-6661/ ISSN(P) 2518-0177
DOI: 10.20448/wsr.v11i1.6223

Caring in AI: Considering the LIDA model

Suereth, Russell

Salve Regina University, Newport, RI, 02840-4192, USA.
Email: russell.suereth@salve.edu

Abstract

The purpose of this article is to consider whether caring can be designed into an artificial intelligence system. Caring is complex. In our daily lives, caring takes several forms and occurs in various ways. This article discusses human caring and how caring has been a vital aspect of our lives since our early ancestors. The research focuses on the form of care a healthcare worker may provide in a clinic or hospital. The research considers whether this attentive form of caring can be designed into AI systems. The approach of this research is to consider a specific AI design known as the LIDA model. The research describes the cognitive cycle and the global workspace within the LIDA model. It also depicts elements in the LIDA model that can be associated with caring. The findings show that caring can occur through gestures and movements. The findings also show that a LIDA agent can perform such gestures and movements and offer an appearance of caring. The findings suggest that a LIDA agent, configured in a particular way, could be a carer in some caring situations. The practical benefit of this research is to show that the LIDA model can be a starting point for designing care in AI systems. Through this research, we may uncover elements of caring that already exist within the LIDA model and can be employed in a caring AI agent.

Keywords: AGI, Artificial intelligence, Caring, Compassion, Emotion, LIDA, NAL, NARS.

Citation | Russell, S. (2024). Caring in AI: Considering the LIDA model. *World Scientific Research*, *11*(1), 6–11. 10.20448/wsr.v11i1.6223
History:
Received: 4 November 2024
Revised: 9 December 2024
Accepted: 16 December 2024
Published: 19 December 2024

Publisher: Asian Online Journal Publishing Group

Funding: This study received no specific financial support.
Institutional Review Board Statement: Not applicable.
Transparency: The author confirms that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.
Competing Interests: The author declares that there are no conflicts of interests regarding the publication of this paper.

Contents

Contribution of this paper to the literature
This article contributes to the literature by considering whether caring aspects can be designed into AI systems. The article focuses on a caring that concentrates on gestures and movements rather than complex realms of human caring. The article considers the LIDA system as a guide for caring in some situations.

1. Introduction

This article addresses whether caring can be designed into an artificial intelligence system. In our everyday lives, caring takes several forms and occurs in varying degrees. The caring a mother has for her baby is perhaps the greatest form of caring possible. However, this high degree of caring is quite challenging to design into an AI system. For example, much occurs in the mother and child space that is hard to explain, or that is even understandable.

Instead, this article focuses on a form of caring that can be easily described in words. This form of caring is the type that a healthcare worker provides in a clinic or hospital. In short, this form of caring is attending to patients who are ill or injured.

This research aims to consider whether care can be designed into AI systems. However, the scope of AI systems is broad. Accordingly, this article reduces the scope by focusing on the LIDA (Learning Intelligent Decision Agent) model. Within this scope, the objectives of this research are:

1) Examine the concept of caring in humans.
2) Investigate the actions that occur in caring.
3) Explore the LIDA model as a foundation for designing care in AI.
4) Identify processes and components of the LIDA model that pertain to caring.
5) Consider how an artificial intelligence agent based on the LIDA model can be a caring agent.
6) Consider the role of a LIDA agent in caring situations.

The research considers literature in the context of caring, artificial general intelligence, and the LIDA model. Accordingly, it employs literature regarding these areas and how caring could be incorporated into a LIDA agent. This research hopes to show that the LIDA framework can be a foundation for designing care in AI systems. It hopes to determine whether the LIDA architecture is flexible enough to add and modify components for specific functions such as caring. The research also hopes to determine whether the LIDA model already has some rudimentary processes and components for the operation of caring in care situations.

2. Human Caring

Since the beginning of our human existence, caring has been an essential component of our lives. Our early human ancestors carried young children in ways that were safe and careful. In that close human connection, other forms of caring may have arisen and continued in later years.

Today, caring is necessary for our development during childhood, for our growth as humans, and in our behavior toward others. We are caring beings, highlighted in the care of our children and our larger family of grandchildren and grandparents. We are also fragile and vulnerable beings; accordingly, care is needed in our infancy and throughout our lives [1]. In a way, simply being aware of the significance of caring is valuable. When we realize its significance, we can see how caring is a crucial component of everything we do.

In our modern world, we care in many different ways. In our health institutions, we look after those who are sick and are suffering [2]. We show caring in our work so that the things we do are done well and in the proper manner. We talk with customers and co-workers in a caring tone because caring is an intrinsic way of how we interact with others and the world around us [3]. It is normal to be careful about the things we do and to be caring for others. Often, that type of caring takes the form of simply paying attention to others. According to Hodges [4] this attention means being careful with them.

This careful attention extends to the things we build and design. That is, we build things in ways that should not harm others, and the things we build should work properly without causing harm. Accordingly, care and safeguards should be designed into the things we build and the processes of building them. As designers of AI systems, we should be aware of the value of care. That is, we should realize the importance of caring in our human lives and work toward designing a type of caring in the AI systems we build. As Ivanhoe, et al. [5] suggest, we should be aware of the importance of care and promote it in our daily lives. This awareness and promotion should also include our work in any of the AI systems we design.

3. Caring as Gestures and Movements

It is sometimes challenging to put a finger on precisely what caring is. That may be because we often regard caring as an emotion that exists inside us. Unfortunately, our internal states, like emotions, are not accessible to other people. In other words, others cannot access what we are actually thinking or feeling [6].

Rather than delving into the hidden internal states of emotions, we could consider a different route. For example, we could consider identifying care by a person's outward gestures and movements. As Meacham and Studley [7] remark, what is important in a caring relationship is not the internal states of those in the relationship but the environment made from gestures and movements that express attentiveness. Interestingly, this reliance on external movements opens the door to accepting the possibility of care relationships between ourselves and robotic carers [7].

Studies show that when carers perform engaging and friendly actions, patient satisfaction is higher. Movements such as sitting, leaning forward, making eye contact, and nodding make patients feel that carers have warmth and concern and that better communication is occurring [8]. However, caring is more than just paying attention to the patient. Caring behaviors also include advocacy, responsiveness, and having a positive attitude [9]. In short, many movements should be considered when designing care in an AI system.

4. Caring and Compassion

Caring is an essential aspect of our social interactions and so is compassion. In the context of this article and the design of care in AI systems, emotions like compassion and sympathy can be viewed as being similar to caring. According to Nussbaum [10] compassion is vital in our everyday lives. For Nussbaum, in our process of providing compassion, we consider how much compassion is deserved in a particular situation and how much we should give [10]. From the perspective of designing AI systems, we can use these considerations as measurements to determine how much compassion an AI system should give.

For Nussbaum [10] caring, compassion, empathy, and sympathy have distinct definitions. However, the differences between these emotions are complex and can be difficult to notice. In designing caring actions in an AI system, we may decide to consider setting aside the nuances of these similar emotions. In other words, it may be prudent to first focus on designing the basic movements of attention, responsiveness, and advocacy in initial designs.

As designers, we should remind ourselves that concepts like compassion and caring are essential aspects of human interaction. For Nussbaum, compassion is required for humans to interact with one another reasonably. Nussbaum [10] doubles down on the requirement by remarking that we would view a person without compassion as a moral disgrace. Nussbaum does not mention AI systems in her book, but we can extrapolate her words to the context of AI systems and how we may feel about an AI machine that is not designed to exhibit caring or compassion.

5. The LIDA Model and Caring

The type of AI described in this article is artificial general intelligence (AGI). This type is general because it is an overall intelligence that is designed for doing many things. In a way, an AGI has similarities to human intelligence since both are robust and multi-capable. However, an actual AGI system is challenging to build — AGI is expected to have a simple form of cognition and cognition is not easy to design or even understand. However, there has been some success in developing AGI systems that can handle various inputs and employ certain principles in selecting actions [11].

This article discusses the LIDA model. LIDA is a comprehensive AGI model with features similar to other AGI systems. A benefit of the LIDA architecture is that it has a modular framework so that different design teams can work on a variety of AI agents that are based on the model [11, 12].

The article focuses on some of the components and processes in LIDA that could be involved in movements of caring. Indeed, much more is going on in the LIDA model than can be described in this article. A more complete description of the system is given in "New Developments in the LIDA Model" [13].

5.1. The Cognitive Cycle in LIDA

In the LIDA model, sensory data depicts the current situation and environment, while high-level principles designate the general types of actions that the AI agent should take. In a sense, the LIDA model is a model of human cognition that uses continuous cognitive cycles of sensation, attention, and action [11]. It is a cycle that mimics the processing humans and animals perform to make decisions [12]. A diagram of the LIDA cognitive cycle is shown in Figure 1.

The awareness of the current situation and the choice of proper actions occur in many different modules within LIDA. For example, the current situation is derived from low-level sensations of the environment and passed to an area of perceptual memory where high-level relations and situations are identified [11]. The action selection process then chooses an action based on the current situation and high-level principles [14].

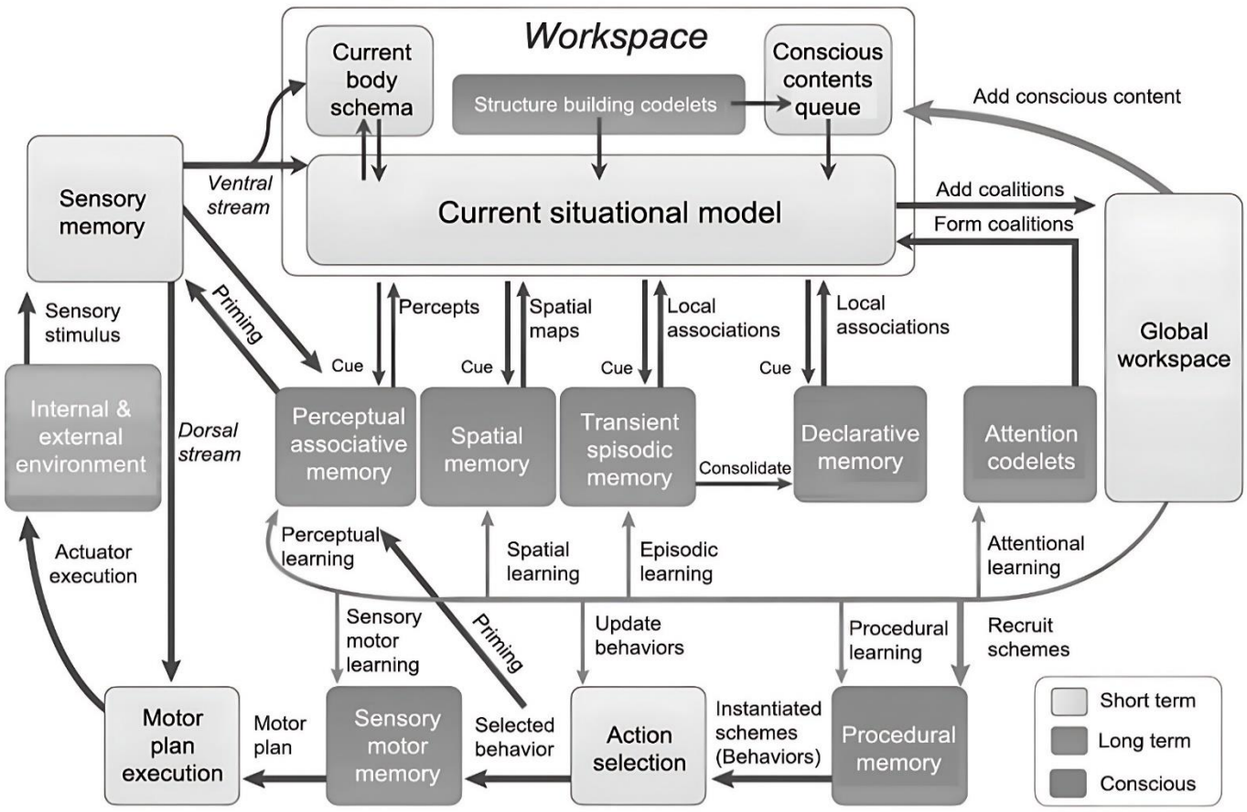

Figure 1. The LIDA cognitive cycle diagram. cognitive computing research group. n.d. Accessed June 10, 2024.
Source: https://ccrg.cs.memphis.edu/index.html.

5.2. The Global Workspace

LIDA employs the idea of a global workspace, which is based on the Global Workspace Theory (GWT) [15]. GWT is a model of how consciousness works. It "postulates that human cognition is implemented by a multitude of relatively small, special purpose processes, almost always unconscious" [16].

GWT conducts the operation of attention by using a competition to determine which situation is the most important at that time. One way of looking at GWT is by viewing it as a spotlight on a theatre stage where the attention is directed to the actor speaking. Where this spotlight shines in GWT is determined by many factors, including sensory inputs such as visual data. The direction of the spotlight is also determined by an assortment of principles that have priority in the attention determination. Such principles may insist that the AI agent not drop drinking glasses or harm humans.

An interesting way of looking at GWT is from Wiersma [17] who suggests that emotions are a generator of consciousness and an initiator of access to the cognitive workspace. Wiersma [17] is a proponent of GWT, which he remarks is valuable because it is simple, and this simplicity helps us use it in different ways. However, for Wiersma, GWT also describes levels of consciousness, as well as levels of emotions. Wiersma's view is in the context of cognitive theory, but the primacy of emotions he describes in GWT could be another way of thinking about caring in the GWT space. In other words, if emotions are a generator of consciousness, then caring also may be a generator of consciousness.

5.3. Caring Elements in LIDA

Within the LIDA framework, several processes and components can be viewed as being associated with caring. In this sense, the LIDA model already contains elements that could enable it to show an outward appearance of caring.

5.3.1. Attention Codelets

From a design viewpoint, codelets are small modules of programming code that serve a specific and somewhat simple function. Many different codelets apply to different situations within the LIDA model. For example, there are intention codelets, expectation codelets, and attention codelets. Attention codelets are designed to focus on certain types of content or situations [18, 19].

In the context of caring, attention codelets can be designed to focus on situations where care occurs. Attention codelets can promote caring by prioritizing situations where caring occurs or is needed. In a waiting room of patients, an AI agent could be configured to attend to the patient who needs the most immediate care. Of course, the AI agent must perform much additional processing to determine which patient needs the most care. Such additional processing would involve sensory input and high-level principles that rank types of injuries and illnesses.

5.3.2. Action Selection

Caring occurs in different ways. One of these is through physical movements that we can observe and are quite evident. Attention, gestures, and responses are some movements that indicate caring moments. In a way, these actions help confirm that caring exists.

Meacham and Studley [7] note that in hospitals and nurseries, an awareness of the internal state of a carer is not necessary. They suggest that the carer's caring actions, or expressive behavior, are important. Accordingly, in these caring environments, caring occurs in movements like attention and gestures.

Caring that transpires through action requires that action occurs. That is, we must have approving facial gestures, a listening demeanor, and an appearance of interest. Ensuring that these movements occur in an AI system requires that these movements are designed into the system. More specifically, it requires designing the selection and execution of those movements and ensuring that they occur at the proper moments.

In the LIDA framework, an AI system's actions are chosen in the action selection process, located at the bottom of the diagram in Figure 1. The choice is based, in part, on the values of LIDA's feelings about that action [20]. The feeling values are a form of motivation that directs the choice of actions [21].

The action selection process could be an advocate for caring by instructing an AI agent to perform caring actions. For example, with some enhancement, the selection process could choose actions that employ medical instruments to check the patient's blood pressure and body temperature. In this context, the selection process would help choose the right caring actions. Additional processing also would be needed to determine when these caring actions should occur and whether they are correctly performed.

5.3.3. Top-Down Principles and Bottom-Up Situations

Performing patient care can be complex. Carers must comply with principles of conduct and adhere to healthcare laws and regulations when caring for patients. However, the actions a carer takes involve more than principles and regulations. Caring is a situational practice that focuses on the patient. The type and degree of care depends on the situation at that moment and with that individual patient. As Poulsen and Burmeister [22] note, carers must make care decisions as the caring practice occurs and while interpreting the patient's needs.

Top-down processes in the LIDA model handle high-level principles and regulations. These top-down processes also handle human-like deliberation and complex decision-making. Bottom-up processes handle sensory data from the current situation. These low-level processes include the details of the latest circumstances and environment. In bottom-up processes, attention is determined by the duration and frequency of the sensory data, which indicates a degree of intensity of situations and feelings [11].

To handle patient care, top-down and bottom-up processes likely need a design enhancement to accommodate caring situations. In this enhancement, bottom-up processes could identify caring situations by focusing on signs of injury or symptoms of illness. In this way, a LIDA agent could detect these signs and symptoms and act accordingly. The action selection process would also need enhancement to choose actions based on several caring aspects, including the type and degree of the injury and symptom. The action chosen could also depend on the agent's assessment of the patient's words, gestures, and reactions. However, the agent's actions are also determined by top-

down processes. In these enhancements, such principles could guide the agent to handle medical instruments and the human patient properly.

Through these enhancements, top-down and bottom-up processes can promote caring by detecting injuries and symptoms, being aware of the current situation, and applying valued principles and traditional methods to attend to patients. From a design viewpoint, the LIDA model appears already designed with processes and components needed to act in a caring manner.

5.3.4. Values of Feelings

In our everyday lives, we choose to do certain activities due to how we feel about them. When given the choice, we tend to do activities that make us feel good. The LIDA model does not have feelings, but it does have some elementary feeling processes. In these processes, ranges of values designate low or high degrees of those simulated feelings. The value of a feeling can determine the focus of attention in a situation or the action that is chosen [23].

In a way, the values of feelings in the LIDA model describe what the LIDA agent likes and dislikes. In other words, the values describe the agent's value system [24]. However, the agent cannot pay attention to only the things that make it happy or do only those activities that make it feel good. The agent also has principles that guide its processing and decision-making.

However, the handling of principles in the LIDA model should be enhanced to become more powerful. Such an enhanced design could focus on abstract and detailed levels of principles so that the principles could apply to various caring situations.

6. LIDA, The Caring Assistant

Caring for others is not easy. It is often a thankless job, and the communication challenges between carers and patients make it more challenging. Leavitt and Leavitt [8] note that successful diagnosis and treatment depend on productive communication between the carer and the patient. However, productive communication is challenging. The authors present examples that highlight the challenge: carers and patients disagree on the main health problem more than 25 percent of the time, about 50 percent of medical problems are not discussed, and carers give full medication instructions less than 60 percent of the time. Unfortunately, these examples are only a sample of the miscommunication that the authors depict.

The challenge of caring is compounded by compassion fatigue. This exhaustion arises from the overload of support needed for patients and their families [25]. The fatigue results in burnout and traumatic stress [26].

One way to resolve some of these caring challenges may be through a caring AI agent. For example, AI systems do not succumb to tiredness or forgetfulness. AI systems would immediately have a patient's complete history and current medications. AI systems can communicate endlessly, all day and night, without getting tired or needing to get home for a family birthday party. AI systems can listen to the same story without being bothered. Furthermore, AI systems can experience the emotions of the patient's family all day without a change in performance. Indeed, designing an AI carer that replaces a human carer is daunting work, but perhaps an AI carer could specialize in specific caring tasks.

6.1. Patient-Centered Care

Patient-centered care, of course, is a focus on the patient. In this caring perspective, the carer explores the patient's experience of illness or injury and centers on understanding the patient as a whole person [8]. Here, the discussions between the carer and patient emphasize the patient's background, family, environment, and how the illness affects the patient's life.

These detailed discussions may be a good task for an AI carer that asks questions and gathers information. Such AI agents could be adept at immediately looking up information for verification and providing further insights in the patient discussion. According to Leavitt and Leavitt [8] patients like patient-centered care and it is effective. That is, patient-centered care is connected to improved health, increased care efficiency, and better adherence to medication regimens.

6.2. Patient-Centered Care in LIDA

The LIDA model is an advanced architecture for artificial general intelligence. However, its capabilities must be enhanced to accommodate patient-centered care. For example, values must be set to specify particular feelings and actions. Such values would be based on a particular type of patient or individual patient. Values could also be assigned to certain principles and caring behaviors specific to a particular patient [22]. We could extend this personalized AI care by considering the specific situation so that care occurs in a manner that is suitable for that patient at that moment [22].

Within the LIDA framework, further customized care could exist through a recalibration. In other words, the needs of a particular patient would likely change through the duration of the patient's journey of injury, hospital stay, and rehabilitation. Accordingly, the patient's environment would also change. A recalibration based on these external situations and resetting the various LIDA values of feelings, actions, and principles can better accommodate the patient's changing health and surroundings [22].

It is interesting to think of a LIDA agent as a care assistant. The role makes sense when we consider that a team of carers provides a variety of caring capabilities and experiences. According to the National Academies of Sciences Engineering and Medicine [27] an optimal team includes various clinical and nonclinical team members. Indeed, designing many types of caring in an AI system would be challenging. A design focusing on caring communication through question-and-answer sessions may be a more reasonable role for a young AI carer.

7. Conclusion

This article began by describing caring among humans. It showed how internal states of caring are not visible and how external caring, the movements of caring, are crucial in caring environments. It provided a short overview of the LIDA model and GWT. At a high level, the article described components of the LIDA model that are or could

be associated with caring. The article discussed patient-centered care and how this form of caring could be performed by a caring agent based on the LIDA model. It suggested that a caring agent could aid in reducing communication challenges between carers and patients and reduce the impact of compassion fatigue. It suggested that a caring agent based on the LIDA model could be designed as a caring assistant in caring environments.

In the LIDA model, caring capabilities could be handled in various ways and locations.

These multiple processes and locations make sense when considering caring in our own human selves. We do not give or receive care in only one particular way or one type of situation. The activities of care happen in our everyday world, where many scenarios and ways of living exist. In a sense, like an enhanced LIDA model, we care in different ways and by using different parts of our human selves.

Based on these findings, the LIDA model already contains components that can be valuable for a caring AI system, although enhancements are needed. Within the LIDA model, attention codelets could direct an AI agent to focus on situations where caring occurs. Action selection in LIDA could direct an AI agent to choose caring actions in caring situations. Top-down processes could direct an AI agent to act according to principles of conduct and ethical rules, while bottom-up processes could detect signs of injury and symptoms of illness. Values assigned to simulations of feelings within the LIDA model could help determine an AI carer's course of actions.

The LIDA model already has many processes and components that can be employed in a caring environment. However, more design work is needed to determine how a LIDA model can be configured for caring. Furthermore, this design work may reveal aspects of the model that need enhancement for caring situations. One possible enhancement may be the design of additional components that handle specific modes of caring and types of patients. Another direction for future design work could focus on a different AI model. One interesting model is the Non-Axiomatic Reasoning System (NARS), which employs a Non-Axiomatic Logic (NAL) to perform reasoning [28]. NARS is a recent AGI architecture that could be used to develop a variety of intelligent agents. An investigation of NARS may show that it can be valuable in developing agents that are employed in a caring environment.

References

[1] L. Mortari, *Care: The primacy of being. In Care Ethics in the Age of Precarity.* Minneapolis, MN: University of Minnesota Press, 2021.

[2] B. Waldenfels, *Care of the self and care of the other. In Care Ethics and Phenomenology: A Contested Kinship. F. Vosman, and P. Nortvedt, Eds.* Bristol, CT: Peeters Publishers, 2020.

[3] A. Alacovska, "From passion to compassion: A caring inquiry into creative work as socially engaged art," *Sociology*, vol. 54, no. 4, pp. 727-744, 2020. https://doi.org/10.1177/0038038520904716

[4] B. H. Hodges, "Carrying, caring, and conversing: Constraints on the emergence of cooperation, conformity, and language," *Interaction Studies: Social Behavior and Communication in Biological and Artificial Systems* vol. 18, no. 1, pp. 26-54, 2017. https://doi.org/10.1075/is.18.1.02hod

[5] P. Ivanhoe, O. Flanagan, V. Harrison, E. Schwitzgebel, and H. Sarkissian, *The oneness hypothesis: Beyond the boundary of self.* New York: Columbia University Press, 2018.

[6] M. Coeckelbergh, "The moral standing of machines: Towards a relational and non-Cartesian moral hermeneutics," *Philosophy & Technology*, vol. 27, pp. 61-77, 2014. https://doi.org/10.1007/s13347-013-0133-8

[7] D. Meacham and M. Studley, *Could a robot care? It's all in the movement. In Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence. P. Lin, K. Abney, and R. Jenkins, Eds.* New York: Oxford University Press, 2017.

[8] J. Leavitt and F. Leavitt, *Improving medical outcomes: The psychology of doctor-patient visits.* Lanham, MD: Rowman and Littlefield, 2011.

[9] C. L. Montgomery, *Healing through communication: The practice of caring.* Thousand Oaks, CA: Sage, 1993.

[10] M. C. Nussbaum, *Upheavals of thought: The intelligence of emotions.* New York: Cambridge University Press, 2003.

[11] W. Wallach, S. Franklin, and C. Allen, "A conceptual and computational model of moral decision making in human and artificial agents," *Topics in Cognitive Science*, vol. 2, no. 3, pp. 454-485, 2010. https://doi.org/10.1111/j.1756-8765.2010.01095.x

[12] J. Snaider, R. McCall, and S. Franklin, "The LIDA framework as a general tool for AGI," presented at the Artificial General Intelligence: 4th International Conference, AGI 2011. Springer Berlin Heidelberg, 2011.

[13] S. Franklin and Cognitive Computing Research Group, *The mind according to lida–a brief account. S. Franklin and the Cognitive Computing Research Group of the University of Memphis, at corg. CS. Memphis. edu.* Memphis, TN, 2018.

[14] S. Franklin, *A foundational architecture for artificial general intelligence. In Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms. B. Goertzel, and P. Wang, Eds.* Amsterdam, NL: IOS Press, 2007.

[15] S. Franklin, "Building life-like 'conscious' software agents," *AI Communications*, vol. 13, no. 3, pp. 183-193, 2000.

[16] B. J. Baars and S. Franklin, "An architectural model of conscious and unconscious brain functions: Global workspace theory and IDA," *Neural Networks*, vol. 20, no. 9, pp. 955–961, 2007.

[17] E. J. Wiersma, "How sustainable are different levels of consciousness?," *The Journal of Mind and Behavior*, vol. 39, no. 3, pp. 155-180, 2018. https://doi.org/10.31234/osf.io/e6497

[18] U. Faghihi and S. Franklin, *The LIDA model as a foundational architecture for AGI. In Theoretical Foundations of Artificial General Intelligence. P. Wang, and B. Goertze, Eds.* Amsterdam, NL: Atlantis Press, 2012.

[19] S. Kugele and S. Franklin, "Learning in LIDA," *Cognitive Systems Research*, vol. 66, pp. 176-200, 2021.

[20] S. Franklin, S. Strain, R. McCall, and B. Baars, "Conceptual commitments of the LIDA model of cognition," *Journal of Artificial General Intelligence*, vol. 4, no. 2, p. 1, 2013. https://doi.org/10.2478/jagi-2013-0002

[21] S. Franklin, S. Strain, S. Kugele, T. Madl, N. Khayi, and K. Ryan, "New developments in the LIDA model," *American Philosophical Association Newsletters: Philosophy and Computers*, vol. 17, no. 2, pp. 10–14, 2018.

[22] A. Poulsen and O. K. Burmeister, "Overcoming carer shortages with care robots: Dynamic value trade-offs in run-time," *Australasian Journal of Information Systems*, vol. 23, 2019. https://doi.org/10.3127/ajis.v23i0.1688

[23] S. Franklin, T. Madl, S. D'mello, and J. Snaider, "LIDA: A systems-level architecture for cognition, emotion, and learning," *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 1, pp. 19-41, 2013. https://doi.org/10.1109/tamd.2013.2277589

[24] S. Franklin and F. Patterson, "The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent," in *IDPT-2006 Proceedings (Integrated Design and Process Technology): Society for Design and Process Science, San Diego, CA*, 2006.

[25] M. Perregrini, "Combating compassion fatigue," *Nursing*, vol. 49, no. 2, pp. 50-54, 2019.

[26] D. Barron, "Burned out on burnout? Strategies to combat compassion fatigue in nursing," *Oklahoma Nurse*, vol. 68, no. 4, pp. 10-12, 2023.

[27] National Academies of Sciences Engineering and Medicine, *Implementing high-quality primary care: Rebuilding the foundation of health care.* Washington, DC: The National Academies Press. https://doi.org/10.17226/25983, 2021.

[28] P. Wang, *Non-axiomatic logic: A model of intelligent reasoning.* Hackensack, NJ: World Scientific, 2013.