Some Research Directions for Multi-Agent Systems in Medicine and Health

Vincenzo Auletta[1], Francesco Cauteruccio[1], Diodato Ferraioli[1,*] and Grazia Ferrara[1]

[1]*Università degli Studi di Salerno, via Giovanni Paolo II 132, Fisciano (SA), 84084, Italy*

Abstract

In this work, we present some directions for research in multi-agent systems that may have useful and interesting applications in medicine and healthcare. Specifically, we focus on three possible applications: improving our understanding of how complex behavior may arise from the interaction of brain cells and how this can be affected; improving our understanding of metastasis formation, and how this can be inhibited; and improving the diffusion of preventive healthcare programs, particularly among hard-to-reach and more vulnerable segments of the population.

1. Introduction

In the last years we observed an explosion of AI tools applied to medicine and health. Examples can be found in advanced imaging analysis [1], predictive analysis [2], robotic surgery [3], and personalized medicine [4]. Most of these tools exploit recent developments in modeling agents and enhancing their learning capacities. However, a marginal role in this explosion has been played by the area concerning the analysis and control of multi-agent systems (MAS) [5]. In this area, the goal is to investigate the global behavior of a complex system arising from the communication, coordination, and interaction among (non necessarily collaborating) multiple agents, and to engineer tools able to drive this behavior towards desired goals.

MASes recently achieved successful applications in many different settings ranging from the design of AI-based agents able to super-human performances in poker and other similar games [6], to the successful organization of patrolling in airports, ports, and reserves (see, e.g., [7]).

We believe that the area of MAS may provide an important contribution to medicine and health fields. Indeed, in this work, we are willing to enlist and discuss a few of the possible research areas in medicine and health that may benefit from the adoption of MAS tools. Note that neither this list, nor the specific problems discussed therein, has to be thought to be exhaustive: we are just focusing on a few problems/directions about which we have the expertise to contribute. Still, we hope that this list can provide inspiration and useful references to AI researchers to contribute both to the progress of MAS to make them more tailored to medicine and health applications, and to the progress of medicine and health through the adoption of MAS tools.

2. Multi-Agent Systems for Modeling the Brain Structure

The brain is a very complex systems that consists of billions of cells connected by synapses. Only recently we were able to fully reconstruct the entire brain connections for some simple animal organisms [8], and only partial connections for humans [9].

Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy
*Corresponding author.
auletta@unisa.it (V. Auletta); fcauteruccio@unisa.it (F. Cauteruccio); dferraioli@unisa.it (D. Ferraioli); grferrara@unisa.it (G. Ferrara)
https://docenti.unisa.it/vincenzo.auletta (V. Auletta); https://docenti.unisa.it/francesco.cauteruccio (F. Cauteruccio); https://docenti.unisa.it/diodato.ferraioli (D. Ferraioli)
0000-0002-7875-3366 (V. Auletta); 0000-0001-8400-1083 (F. Cauteruccio); 0000-0002-7962-5200 (D. Ferraioli); 0009-0005-1691-5672 (G. Ferrara)

It has been evident for a long time that the complex behaviors and functions of the brain (perception, cognition, memory, emotion, etc) cannot be understood by the individual study of brain cells and/or of their synapses. These behaviors and functions are a result of the interaction among different brain cells and or synapses. This has naturally lead researchers to model the brain through networks, where nodes (i.e., neurons) are connected through links (synapses, gap junctions), and to use tools of network science, such as the analysis of degree distribution, community structure, and degree distribution, to improve the analysis of both the structure and the functionality of the brain as well as its evolution [10]. The network model of the brain and the tools from network science are still insufficient to capture most of the complex behaviors of this organ. This is because a network representation of the brain still fails to model important features. For example, the connections among neurons may happen through synapses, gap junctions, neuromodulators, and circulating gut peptides, with each different connection having different properties [11]. Similarly, brain cells also have different functional and structural properties, since they change in geometry and morphology, and hence signals received through connections may affect their metabolism in different ways [12].

Anyway, MAS and Network Science have been able since long time in modeling complex systems and studying their structural properties that enable global behaviors from the interaction of their components [5]. For example, in MAS, researchers have considered more complex interactions than simple dyadic links, as, for example, interactions about coalitions of agents, possibly hierarchically organized. This leads to a representation of the brain through more complex structures such as multi-layer [13] and hypergraph networks [14]. These would help to model the difference in connections among brain cells. Research in MAS, moreover, may help to find useful structural properties of these complex networks that would be useful for improving our understanding of the brain (see, e.g., the concept of segregation [15] and its modeling of how specialized subsystems of the brain maintain distinct roles while still integrating within global functions). Similarly, research in MAS has been able to model interactions between different agents, differing both in their type and in how events affect their welfare. This can be achieved through game-theoretic models and techniques. This kind of research allows for the prediction of emerging behavior from the interaction among these different agents, through well-known concepts of Nash equilibria or more fine-grained stability concepts [16, 17].

Another relevant feature of brain structure that has been often observed is its dynamical nature [18]: interactions among brain cells evolve over time, with old connections disappearing and new connections being created. Even more interesting, this structural dynamics co-evolves with the evolution of functionalities [19]: some functionality allows synapses to form and change, but this in turn enables new functionalities. Again, these dynamical and co-evolving aspects of interactions are not new to MAS research (see, e.g., [20]).

Hence, it may be beneficial to deepen this research, trying to embed in these models specific features of brain cells and their interaction, and for the other side to use these tools for explaining and predicting complex behavior of the brain, or also to engineer possible interventions.

3. Multi-Agent Systems and Metastatic Formation Understanding

Metastasis [21] is a complex process where cancer cells spread from the primary tumor to distant organs in the body. While significant advances have been made in treating localized primary tumors through surgery, radiation, and targeted therapies, metastatic disease remains largely incurable and accounts for approximately 90% of cancer-related mortality [22]. This stark statistic underscores the critical need for improved understanding of metastatic processes and the development of more effective therapeutic strategies targeting metastasis.

The clinical management of metastatic disease presents unique challenges compared to primary tumors. While primary tumors typically occupy a single anatomical location and may be amenable to direct intervention, metastatic disease is characterized by multiple tumor deposits dimetastatic cells often exhibit enhanced resistancestributed across diverse organ systems [23]. This spatial distribution not only complicates surgical approaches but also necessitates systemic treatments that must navigate complex barriers while minimizing toxicity to healthy tissues [24]. Furthermore, metastatic cells often exhibit enhanced resistance to conventional therapies, having already survived the selective pressures of the metastatic cascade, including detachment from the primary tumor, survival in circulation, and colonization of foreign tissue microenvironments. A particularly intriguing aspect of metastasis is its non-random nature. Specific cancer types exhibit predictable patterns of organ-specific metastasis, a phenomenon termed "metastatic tropism" [25]. This and other similar patterns suggest that metastasis is not simply a passive stochastic process. Understanding the mechanisms driving these organ preferences could reveal critical vulnerabilities in the metastatic process and identify novel therapeutic targets.

Today, metastasis is understood as a complex, multi-step process involving bidirectional communication between tumor cells and host microenvironments, with roles for primary tumors, circulating tumor cells, distant organs, and the immune system [23]. This dynamic interplay suggests that metastatic cells engage in sophisticated decision-making processes when selecting and adapting to new tissue environments; a perspective that invites modeling approaches capable of capturing these strategic interactions.

Research in MAS faced the interaction among this strategic components several times: this has been concretely and efficiently modeled through game-theoretic and computational tools. Various game-theoretic approaches have been already applied in cancer research to understand different mechanisms including tumor growth, invasion dynamics, cooperation among cancer cells and treatment resistance [26, 27, 28]. Still these contributions mainly focus on the evolution of the metastasis within an organ and its treatment resistance, but they fail to explain other features, such as metastatic tropism, i.e., why some metastatic cells happens to aggregate mainly in specific organ microenvironments, and not in others. We believe that MAS reasearch can provide useful tools to address this issue. For example, hedonic games [29] have been heavily used for modeling how agents form stable coalitions based on their preferences for being part of different groups. This parallels the biological situation where metastatic cells “choose” specific organ microenvironments that maximize their survival and proliferation potential.

4. Multi-Agent Systems for Public Health Prevention Campaigns

Prevention is fundamental in order to guarantee individual health. This is even more important for public health, i.e., protecting the population from infectious diseases. While this aspect has been recently highlighted by the COVID-19 pandemic, it emerges as relevant in most of the medical literature about these diseases. For example, the 2024 report of the Joint United Nations Programme on HIV/AIDS (UNAIDS) (https://crossroads.unaids.org/) states that “much more effort and urgency is required to accelerate prevention”, especially towards “marginalized people” to address the goal “to end AIDS as a public health threat by 2030”. Can AI help in identifying critical subjects that can help in enlarging the participation to prevention campaigns, in particular among categories of marginalized people that are difficult to reach with the normal information channels (e.g., homeless, sex-workers, drug-addicted, prisoners), i.e., the ones for whom the access to campaign is usually more difficult? Can AI help to "convince" these people to enroll in these prevention campaigns, and prevent them from dropping out?

We believe that research in MAS has the potential to contribute in this direction. Indeed, the study of influence maximization and opinion formation [30] has recently become a very active subject within this area, where the objective is to analyze how the information diffuses in a network, how the formation of opinion is affected from this diffusion, and how it is possible to influence these processes. Some of these tools have already been successfully applied to public health prevention. For example, some studies have approached the problem of HIV prevention initiatives among homeless youth from the perspective of maximizing influence [31], or the prevention of maternal and neonatal health issues among pregnant women [32].

These applications can help research in MAS to identify new problems of interest, whose solution may help to improve the efficiency of health prevention policies. An example of a possible direction is to work with only partial knowledge of the setting, such as partial knowledge of relationships, their effectiveness in spreading information, or prior preferences by people. Online learning techniques have been recently applied to deal with this partial information [33]. It would be interesting to investigate the extent to which these tools can be useful. Other possible direction is to consider (co)-evolving networks, in which the relationship among people depends also on the actions that they take, or privacy-aware agents, to guarantee that revealed health-relevant information cannot damage agents itself.

5. Conclusions

In this work, we discussed some possible research directions in MAS that can have applications in medicine and health. We believe that advances in MAS can be very useful for improving the understanding of living bodies and thus leading to progress in medicine and public health. We invite all interested readers to join us in this research direction.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] H.-P. Chan, R. K. Samala, L. M. Hadjiiski, C. Zhou, Deep learning in medical image analysis, Deep learning in medical image analysis: challenges and applications (2020) 3–21.

[2] R. Fakoor, F. Ladhak, A. Nazi, M. Huber, Using deep learning to enhance cancer diagnosis and classification, in: ICML, volume 28, ACM New York, NY, USA, 2013, pp. 3937–3949.

[3] M. Iftikhar, M. Saqib, M. Zareen, H. Mumtaz, Artificial intelligence: revolutionizing robotic surgery, Annals of Medicine and Surgery 86 (2024) 5401–5409.

[4] N. J. Schork, Artificial intelligence and personalized medicine, Precision medicine in Cancer therapy (2019) 265–283.

[5] G. Weiss (Ed.), Multiagent Systems. 2nd edition, MIT Press, 2013.

[6] N. Brown, T. Sandholm, S. Machine, Libratus: The superhuman ai for no-limit poker., in: IJCAI, 2017, pp. 5226–5228.

[7] J. Pita, M. Jain, F. Ordónez, C. Portway, M. Tambe, C. Western, P. Paruchuri, S. Kraus, Using game theory for los angeles airport security, AI magazine 30 (2009) 43–43.

[8] C. Verasztó, S. Jasek, M. Gühmann, R. Shahidi, N. Ueda, J. D. Beard, S. Mendes, K. Heinz, L. A. Bezares-Calderón, E. Williams, et al., Whole-animal connectome and cell-type complement of the three-segmented platynereis dumerilii larva, BioRxiv (2020) 2020–08.

[9] A. Shapson-Coe, M. Januszewski, D. R. Berger, A. Pope, Y. Wu, T. Blakely, R. L. Schalek, P. H. Li, S. Wang, J. Maitin-Shepard, et al., A connectomic study of a petascale fragment of human cerebral cortex, BioRxiv (2021) 2021–05.

[10] R. F. Betzel, D. S. Bassett, Multi-scale brain networks, Neuroimage 160 (2017) 73–83.

[11] C. Presigny, F. De Vico Fallani, Colloquium: Multiscale modeling of brain network organization, Reviews of Modern Physics 94 (2022) 031002.

[12] E. Bullmore, O. Sporns, The economy of brain network organization, Nature reviews neuroscience 13 (2012) 336–349.

[13] S. Boccaletti, G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, M. Zanin, The structure and dynamics of multilayer networks, Physics reports 544 (2014) 1–122.

[14] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, G. Petri, Networks beyond pairwise interactions: Structure and dynamics, Physics Reports 874 (2020) 1–92.

[15] A. Failla, G. Rossetti, F. Cauteruccio, Beyond boundaries: Capturing social segregation on hyper-networks, in: ASONAM, Springer, 2024, pp. 40–55.

[16] V. Auletta, D. Ferraioli, F. Pasquale, G. Persiano, Metastability of logit dynamics for coordination games, Algorithmica 80 (2018) 3078–3131.
[17] D. Ferraioli, C. Ventre, Metastability of the logit dynamics for asymptotically well-behaved potential games, ACM Trans. Algorithms 15 (2019) 27:1–27:42.
[18] S. Chung, L. F. Abbott, Neural population geometry: An approach for understanding biological and artificial neural networks, Current opinion in neurobiology 70 (2021) 137–144.
[19] C. H. Papadimitriou, S. S. Vempala, D. Mitropolsky, M. Collins, W. Maass, Brain computation by assemblies of neurons, Proceedings of the National Academy of Sciences 117 (2020) 14464–14472.
[20] V. Auletta, A. Fanelli, D. Ferraioli, Consensus in opinion formation processes in fully evolving environments, in: AAAI, 2019, pp. 6022–6029.
[21] S. Gerstberger, Q. Jiang, K. Ganesh, Metastasis, Cell 186 (2023) 1564–1579.
[22] X. Guan, Cancer metastases: challenges and opportunities, Acta pharmaceutica sinica B 5 (2015) 402–418.
[23] D. F. Quail, J. A. Joyce, Microenvironmental regulation of tumor progression and metastasis, Nature medicine 19 (2013) 1423–1437.
[24] P. S. Steeg, Targeting metastasis, Nature reviews cancer 16 (2016) 201–218.
[25] Y. Li, F. Liu, Q. Cai, L. Deng, Q. Ouyang, X. H.-F. Zhang, J. Zheng, Invasion and metastasis in cancer: molecular insights and therapeutic targets, Signal transduction and targeted therapy 10 (2025) 57.
[26] M. Archetti, K. J. Pienta, Cooperation among cancer cells: applying game theory to cancer, Nature Reviews Cancer 19 (2019) 110–117.
[27] B. Wölfl, H. Te Rietmole, M. Salvioli, A. Kaznatcheev, F. Thuijsman, J. S. Brown, B. Burgering, K. Staňková, The contribution of evolutionary game theory to understanding and treating cancer, Dynamic Games and Applications 12 (2022) 313–342.
[28] C. Morison, M. Fic, T. Marcou, J. Mohamadichamgavi, J. Redondo Antón, G. Sayyar, A. Stein, F. Bastian, H. Krakovská, N. Krishnan, et al., Public goods games in disease evolution and spread, Dynamic Games and Applications (2025) 1–17.
[29] A. Bogomolnaia, M. O. Jackson, The stability of hedonic coalition structures, Games and Economic Behavior 38 (2002) 201–230.
[30] V. Auletta, D. Ferraioli, G. Ferrara, How to mitigate disagreement and polarization in opinion formation processes on social networks, in: SPIRIT, 2024.
[31] B. Wilder, N. Immorlica, E. Rice, M. Tambe, Maximizing Influence in an Unknown Social Network, Proceedings of the AAAI Conference on Artificial Intelligence 32 (2018).
[32] S. Pal, M. Tambe, A. S. Suggala, K. Shanmugam, A. Taneja, Improving mobile maternal and child health care programs: Collaborative bandits for time slot selection, in: AAMAS, 2024, pp. 1540–1548.
[33] V. Auletta, F. Carbone, D. Ferraioli, C. Vinci, Adaptive multi-round influence maximization with limited information, in: AAMAS, 2025.