We've Been Putting People to Sleep for 175 Years. And We Still Don't Fully Know How

A Commentary on the Mechanistic Gap in General Anesthesia

Cela fait 180 ans qu'on endort des patients. Et on ne comprend toujours pas complètement comment cela fonctionne.
Commentaire sur le manque d'explication mécanistique en anesthésie générale

Ayman Assaaoudi[1*]

1. University of Ottawa, Ottawa, ON, Canada
*Corresponding author. Email: aassa087@uottawa.ca

Abstract | Résumé

General anesthesia is one of modern medicine's most impressive successes. Every day, patients are made unconscious, kept still and pain-free during surgery, and then brought back to awareness with no memory of the operation itself. Clinically, this works remarkably well. The harder question is why it works. We can describe many of the drugs, receptors, and brain patterns involved, but the explanation is still incomplete when we try to connect a molecular drug effect to the lived disappearance of awareness.
This commentary follows that gap in three steps. It first looks at the early shift from the Meyer-Overton lipid theory to receptor-based explanations involving targets such as GABA-A and NMDA receptors. It then considers why receptor pharmacology alone does not fully explain the anesthetic state, especially because unconsciousness, immobility, amnesia, and analgesia are partly separable effects. Finally, it turns to the deeper problem: anesthesia removes consciousness, but neuroscience still does not have a complete theory of what consciousness is. The result is not a claim that anesthesia is unsafe or mysterious in every respect, but rather that an important mechanistic gap remains even after major clinical and scientific progress.

———

L'anesthésie générale est l'une des plus grandes réussites de la médecine moderne. Chaque jour, des patients sont rendus inconscients, maintenus immobiles et sans douleur pendant une chirurgie, puis ramenés à l'état d'éveil sans aucun souvenir de l'opération elle-même. Sur le plan clinique, cette pratique fonctionne remarquablement bien. La question plus difficile est de comprendre pourquoi elle fonctionne. Nous pouvons décrire plusieurs médicaments, récepteurs et modèles d'activité cérébrale impliqués, mais l'explication reste incomplète lorsqu'on essaie de relier l'effet moléculaire d'un anesthésique à la disparition vécue de la conscience.
Ce commentaire examine cet écart en trois étapes. Il présente d'abord le passage de l'ancienne théorie lipidique de Meyer-Overton vers des explications centrées sur les récepteurs, comme les récepteurs GABA-A et NMDA. Il explique ensuite pourquoi la pharmacologie des récepteurs, à elle seule, ne suffit pas à expliquer complètement l'état anesthésique, surtout parce que l'inconscience, l'immobilité, l'amnésie et l'analgésie sont des effets qui peuvent être en partie séparés. Enfin, il aborde un problème plus profond : l'anesthésie retire la conscience, mais les neurosciences ne possèdent toujours pas de théorie complète de ce qu'est la conscience. L'idée n'est donc pas de dire que l'anesthésie est dangereuse ou entièrement mystérieuse, mais plutôt de montrer qu'un écart mécanistique important demeure, malgré les grands progrès cliniques et scientifiques.

Keywords: General anesthesia, Meyer-Overton correlation,GABA-A receptors, Consciousness, Thalamocortical networks · Xenon anesthesia, Intraoperative awareness

Introduction

General anesthesia was first used publicly in surgery by the dentist William Morton in 1846, when he demonstrated ether anesthesia at Massachusetts General Hospital in front of a group of skeptical surgeons. He anesthetized a patient long enough for a neck tumor to be removed without pain. Within months, the use of general anesthesia spread internationally. Nearly 180 years later, anesthesia is one of the safest areas of hospital medicine. Every day, anesthesiologists provide anesthesia and pain control for hundreds of millions of procedures worldwide each year. The drugs are better, the monitoring is better, and the safety record is extraordinary. That clinical success is exactly what makes the remaining scientific uncertainty so striking.

Anesthetics work, but explaining how they work is harder than it first appears. It is fairly easy to describe the first step: an anesthetic molecule binds to, blocks, or modulates a target in the nervous system. The difficulty is the next part. Somehow, those molecular changes alter circuits, reshape large-scale brain

communication, and produce a patient who is unconscious, immobile, amnesic, and unable to experience pain. The explanation becomes less secure as we move upward from receptors to circuits to the whole brain. Some of the missing pieces are experimental. Others are conceptual, especially because we still do not have a settled account of what consciousness is. This paper follows that chain level by level, pointing out where the evidence is strong and where the explanation still thins out.

The Lipid Theory: A Great Idea That Wasn't Quite Right

The first systematic attempt to explain anesthesia came from a striking observation. In 1899 and 1901, two pharmacologists working independently, Hans Meyer in Germany and Charles Overton in England, noticed something remarkably consistent across a wide range of chemically different anesthetic compounds (1, 2). If you measured how readily each one dissolved into fat versus water, the ranking almost perfectly predicted their potency. More fat-soluble meant more powerful, nearly every time.

The obvious interpretation was that anesthetics worked by dissolving into the fatty membranes of neurons and disrupting them in a nonspecific physical way, much like a drop of oil changes the surface of water. It was an appealing idea because it gave one simple explanation for many different compounds, from ether to chloroform to alcohol. There were researchers who suspected early on that proteins might matter more than lipids, but the lipid theory had something those alternatives did not yet have: a clean, predictive relationship between fat solubility and anesthetic potency. For that reason, it dominated the field until later experiments made the protein-target explanation harder to ignore.

The cracks began to appear slowly. The most convincing evidence against a purely lipid-based explanation came from studying what are known as stereoisomers. These are molecules with identical chemical formulas and identical fat solubility, but with mirror-image three-dimensional shapes. If anesthesia were simply about dissolving into fatty membranes, mirror-image versions of the same molecule should work equally well. But they do not. Stereoisomers of the same compound can differ meaningfully in their anesthetic potency despite being chemically indistinguishable on paper (3).

This finding carries a clear implication. Lipid membranes are indifferent to molecular shape, but proteins are exquisitely sensitive to it. The fact that shape affected potency pointed strongly toward specific protein structures as the real targets, not the membranes themselves.

What Receptors Can and Can't Explain

When researchers started narrowing their focus to specific proteins, two kept surfacing in the literature. GABA-A receptors are a type of ion channel found in the neuronal membrane. Their role is inhibitory, meaning that when they open up, the neuron has a harder time firing. NMDA receptors work in the other direction, promoting excitation and contributing to memory encoding. Pull one up, the other down, and most anesthetics are essentially playing that game in one form or another.

Propofol, which is probably what most people get before surgery these days, works by pushing heavily on the GABA-A side. It does this by boosting inhibitory signals throughout the brain, gradually suppressing neural activity until the patient is under. Clinicians like it because it works fast and wears off cleanly, which is not something every anesthetic can claim.

Ketamine is a stranger case. It does not touch GABA-A much at all. Instead it blocks NMDA receptors, and what follows looks less like sleep and more like a dissociative state, where normal perception just kind of comes apart. That made it useful in settings where you needed something fast and did not have a full anesthesia setup handy, emergency medicine, field surgery, pediatric procedures. What nobody expected was that psychiatrists would eventually come knocking. At sub-anesthetic doses, ketamine turns out to act on depression surprisingly quickly, sometimes within hours, in patients who had failed multiple other treatments. The NMDA receptor, it seems, is involved in more than memory.

Volatile agents such as isoflurane and sevoflurane cast a wider net, affecting GABA-A, NMDA, and several other targets (4, 5). This receptor-level work has been genuinely important, but it does not solve the whole problem. Part of the gap is biological: we still do not fully know which mechanisms are responsible for which parts of the anesthetic state. Part of it is conceptual: if the endpoint is loss of consciousness, then we also need to know what consciousness is before we can fully explain what has been removed.

What we call "general anesthesia" is not one thing. It is actually four distinct phenomena happening together: unconsciousness, immobility, amnesia, and pain relief. And these four components appear to involve different mechanisms in different parts of the nervous system. The immobility component is primarily a spinal cord effect. The unconsciousness component is cortical. Amnesia involves hippocampal circuits specifically. These are anatomically and mechanistically separate processes.

Why does this matter? Because a drug that only targets spinal mechanisms could produce complete immobility in a fully conscious patient. That patient would be awake, aware, and experiencing everything but not able to move or speak. It is what intraoperative awareness is, and it happens to approximately one or two patients per thousand procedures (11). The fact that we cannot reliably prevent it, despite decades of effort, reflects how incompletely we understand which mechanisms are actually responsible for which components of the anesthetic state.

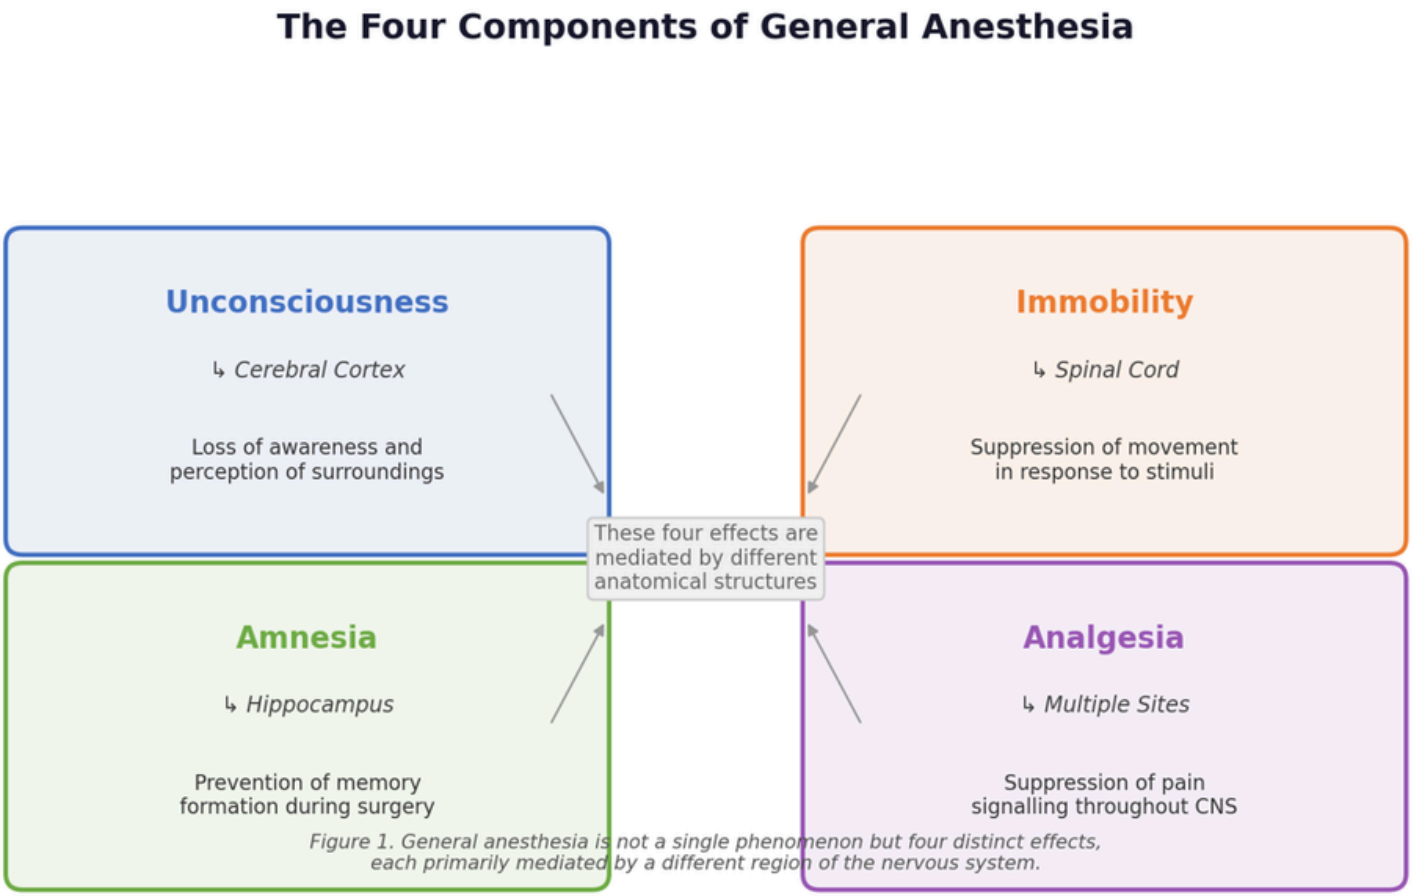

Figure 1. The four components of general anesthesia.

Xenon: The Part That Should Bother Us More

Xenon makes this problem especially clear, although it should not be treated as magic. It is a noble gas, in the same family as helium and neon, and it does not react with biological molecules in the usual chemical sense. It has no functional groups, no charge, and no obvious chemical handle that would make it look like a conventional drug. It is also not metabolised; when it is removed, the gas leaves the body essentially unchanged. Yet at a sufficient concentration, xenon can produce general anesthesia (6).

The best current explanation is still physical rather than chemical in the usual sense. Xenon can occupy hydrophobic pockets in certain receptor proteins and, by sitting there, change how those proteins move or function (7). That explanation is plausible and supported by structural and pharmacological evidence. Still, xenon is useful because it forces the question into sharper focus.

If a chemically inert atom can help produce unconsciousness, then anesthesia cannot be explained only by the familiar picture of drugs forming specific chemical interactions with targets. The mechanism has to be broad enough to include both complex pharmaceutical molecules and a simple noble gas. That does not mean receptor biology is wrong. It means the deeper explanation has to connect physical binding, protein dynamics, neural networks, and consciousness in a way we still have not fully achieved.

The Consciousness Problem Underneath All of This

This is where anesthesia becomes more than a standard pharmacology problem. We have many drugs and gases that reliably produce unconsciousness, but explaining that effect requires knowing what has actually been lost. That is still difficult. We can measure brain activity as consciousness fades, and this has been extremely useful. For example, Emery Brown and colleagues have shown that different anesthetics produce different EEG patterns. Propofol is often associated with slow oscillations around 1 Hz and frontal alpha activity, whereas ketamine is associated with higher-frequency gamma activity that is more posterior (8). These patterns help clinicians and researchers track anesthetic depth, but they are still correlates of unconsciousness, not a complete explanation of consciousness itself.

At a broader anatomical level, one influential account focuses on the thalamus, a deep brain structure that helps relay information between sensory systems and the cortex. The thalamocortical hypothesis proposes that consciousness depends on ongoing communication between the thalamus and the cortex, and that anesthetics disrupt this loop (9). That idea is important, but it is probably not the whole story. Cortical connectivity, the default mode network, and brainstem arousal systems also appear to matter, and different anesthetics may reach unconsciousness through somewhat different network routes.

The same disruptions of thalamocortical interactions can occur during some stages of normal sleep or during seizures, and are not typically followed by a complete loss of consciousness. Thus, thalamocortical disruption may be a necessary component of the underlying causes of coma, but it is not sufficient. And competing theories of consciousness — Integrated Information Theory, Global Workspace Theory, others — make somewhat different predictions about what should happen under anesthesia that the current evidence has not cleanly resolved (10).

Why Does This Matter?

This is not just a theoretical issue. It matters for patient safety, informed consent, and the way anesthetic depth is monitored. Intraoperative awareness is the clearest example. The main monitoring tool used to reduce this risk is the Bispectral Index, or BIS, which turns EEG activity into a score meant to estimate anesthetic depth. BIS can be useful, but it is not perfect. It was developed from observed EEG patterns rather than from a complete biological theory of what anesthetic depth really is. As a result, two patients can have similar BIS values while being in different underlying brain states, especially under different anesthetic drugs.

There is also a practical dosing problem. People vary in how much anesthetic they require, and this variability can be influenced by age, genetics, prior drug exposure, medical history, and baseline neurological function. Some patients need more than the average dose, while others need less. Because this variability cannot yet be predicted from first principles, dosing still relies heavily on population averages, clinical signs, and monitoring tools. That creates room for error in both directions: too little anesthesia risks awareness, while too much can increase physiological stress and delay recovery.

Future work should therefore treat anesthesia as a network-level brain state, not only as a list of receptor effects. High-density EEG and fMRI are now much better suited to studying how anesthetics change communication among brain regions. This kind of work could also help with the long-term goal of separating the four components of anesthesia more precisely: unconsciousness, immobility, amnesia, and analgesia. A drug or monitoring strategy that targets those components separately would require a much clearer map of how each one is produced (12).

Conclusion

Anesthesia is safe and effective, and that achievement is remarkable. Still, nearly 180 years after Morton's ether demonstration, a complete mechanistic account remains out of reach. The field has learned a great deal about receptor targets,

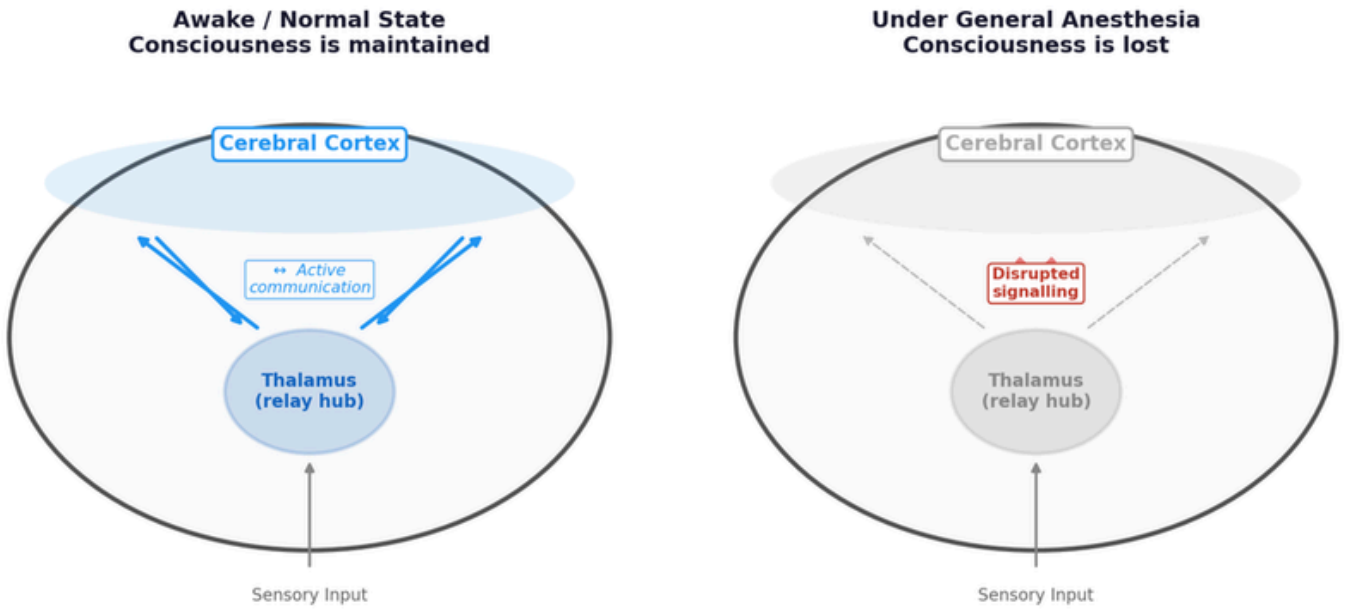

Figure 2. The thalamocortical hypothesis of anesthesia. consciousness depends on constant two-way communication between the thalamus and the cortex (left). Anesthetic drugs disrupt this communication (right), interrupting the neural activity thought to underlie awareness.

EEG signatures, and large-scale brain networks. The difficulty is that these findings do not yet join into a full explanation of how an anesthetic removes consciousness. That question sits between neuroscience, pharmacology, and philosophy, which is partly why it has been so hard to close.

There is a real difference between using anesthesia safely through empirical dosing and monitoring, and understanding it well enough to predict each patient's response from first principles. That gap has practical consequences, including intraoperative awareness, dosing variability, and the future design of anesthetics that could separate the different components of the anesthetic state. The science of anesthesia has come an extraordinary distance since 1846. It simply has not come all the way.

References

1. Meyer, H. H. (1899). Zur Theorie der Alkoholnarkose. Archiv für experimentelle Pathologie und Pharmakologie, 42, 109–118. (No DOI — predates digital era)
2. Overton, C. E. (1901). Studien über die Narkose. Gustav Fischer. (No DOI — book)
3. Franks, N. P., & Lieb, W. R. (1991). Stereospecific effects of inhalational general anesthetic optical isomers on nerve ion channels. Science, 254(5030), 427–430. https://doi.org/10.1126/science.1925602
4. Franks, N. P., & Lieb, W. R. (1994). Molecular and cellular mechanisms of general anaesthesia. Nature, 367(6464), 607–614. https://doi.org/10.1038/367607a0
5. Franks, N. P. (2008). General anaesthesia: From molecular targets to neuronal pathways of sleep and arousal. Nature Reviews Neuroscience, 9(5), 370–386. https://doi.org/10.1038/nrn2372
6. Sanders, R. D., & Maze, M. (2005). Xenon: From stranger to guardian. Current Opinion in Anaesthesiology, 18(4), 405–411. https://doi.org/10.1097/01.aco.0000174957.97759.f6
7. Dickinson, R., Peterson, B. K., Banks, P., Simillis, C., Martin, J. C. A., Valenzuela, C. A., Maze, M., & Franks, N. P. (2007). Competitive inhibition at the glycine site of the N-methyl-D-aspartate receptor by the anesthetics xenon and isoflurane. Anesthesiology, 107(5), 756–767. https://doi.org/10.1097/01.anes.0000287061.77674.71
8. Purdon, P. L., Pierce, E. T., Mukamel, E. A., Prerau, M. J., Walsh, J. L., Wong, K. F. K., Salazar-Gomez, A. F., Harrell, P. G., Sampson, A. L., Cimenser, A., Ching, S., Kopell, N. J., Tavares-Stoeckel, C., Habeeb, K., Merhar, R., & Brown, E. N. (2013). Electroencephalogram signatures of loss and recovery of consciousness from propofol. Proceedings of the National Academy of Sciences, 110(12), E1142–E1151. https://doi.org/10.1073/pnas.1221180110
9. Alkire, M. T., Hudetz, A. G., & Tononi, G. (2008). Consciousness and anesthesia. Science, 322(5903), 876–880. https://doi.org/10.1126/science.1149213
10. Tononi, G. (2008). Consciousness as integrated information: A provisional manifesto. Biological Bulletin, 215(3), 216–242. https://doi.org/10.2307/25470707
11. Mashour, G. A., Orser, B. A., & Avidan, M. S. (2011). Intraoperative awareness: From neurobiology to clinical practice. Anesthesiology, 114(5), 1218–1233. https://doi.org/10.1097/ALN.0b013e31820fc9b6
12. Hemmings, H. C., Jr., Riegelhaupt, P. M., Kelz, M. B., Solt, K., Eckenhoff, R. G., Orser, B. A., & Goldstein, P. A. (2019). Towards a comprehensive understanding of anesthetic mechanisms of action: A decade of discovery. Trends in Pharmacological Sciences, 40(7), 464–481. https://doi.org/10.1016/j.tips.2019.05.001