

Between Flesh and Code: Free Translations from Biology to Computer Algorithms

ABSTRACT

Biologists have long criticized the oversimplification in the biologically-rooted analogies in computer science. In the development of modern neural networks, noting that real neurons are vastly more complex than artificial nodes, and that brains do not learn via backpropagation the way our algorithms do. Instead of seeing this disparity as a failure, I argue it is the point: like a translator capturing a poem's essence over exact wording, bio-inspired algorithms carries the original meaning across when they embrace interpretation over imitation. This talk emerges from my experience in translator and from my occasional encounters with biologically inspired algorithms during my study in computer science. Repurposing parameters in translation theory as the meta-framework, the mapping between biological mechanism and computer algorithm are dissected in dialogue. In particular, five central ideas, supermemes, from translation theory, is utilized to examine how knowledge moves and transforms between biology and computing: source, target, equivalence, untranslatability, free versus literal, and the notion that all writing is translation. By representing the exchange between organism and machine as a translation, a new vision may be conjured, one that understand bio-inspired algorithms as free translations across epistemic boundaries.

KEYWORDS bio-inspired computing, artificial intelligence, reinforcement learning, behavioral biology, translation theory, untranslatability

INTRODUCTION

In the months leading up to graduate study in computer science, my reading was split between primary literature in artificial intelligence and foundational texts in translation theory, prompted by my translation of a contemporary work on bioart ethics, *Human Germline Gene Editing is Bioart: An Open Letter to Lulu and Nana* by Dr. Adam Zaretsky. The juxtaposition was initially jarring, but it soon made a familiar pattern hard to ignore. That pairing of translation practice on one side and algorithmic literature on the other reframed bio-inspired computing as a problem of transfer rather than imitation.

Many of the tools that now feel native to the target-domain of computer science were first articulated as interpretations of the source-domain of biology. Early computer vision offers a classic example: long before “feature extraction” became a standard engineering phrase, neurophysiology had already shown that the mammal visual systems contain units that respond selectively to oriented edges. Computer scientists later built edge and orientation detectors into algorithms and architectures that treat those primitives as the beginning of seeing. Evolutionary algorithms provide another case, where ideas about variation, selection, and population dynamics are recast into search procedures that mutate candidate solutions and retain “winners” as a basis for the next generation. Beyond these well-worn examples, bio-inspired computing is filled with quieter translations. Artificial immune system algorithms borrow from clonal selection, negative selection, and immune memory to frame anomaly detection and adaptive classification. Swarm-based optimization methods translate collective behavior, such as ant foraging and pheromone trails, or flocking dynamics, into distributed search and coordination. Even microbial life has become a source text. Bacterial foraging optimization and related heuristics treat chemotaxis-like movement as a way to explore and exploit a landscape of possibilities. Each of these examples begins with a biological mechanism, then performs a conceptual rewrite so that it can function inside the mathematics and infrastructure of computing.

It became hard not to see a family resemblance. When computer scientists look to biological mechanisms for inspiration, they are, in effect, translating. The source-domain of biology is rendered into the target-domain of computer science. Something is preserved, something is lost, and something genuinely new appears in the crossing. This essay elaborates that intuition and offers a conceptual framework drawn from translation theory for evaluating and defending bio-inspired algorithms, without conflating them with the living systems that inspire them. The discussion expands the arguments I presented at TTT 2025.

I begin with a compact case study in reinforcement learning (RL) because it offers a rare, traceable loop of translation between domains. Concepts from animal learning and behavioral psychology are translated into the computer

science domain as RL algorithms, and those algorithms later return as interpretive frameworks for neurobiological data on learning signals (Dayan, 1992; Kaelbling, Littman, & Moore, 1996). In the case of RL, the source-domain does not simply donate metaphors to the target-domain; the target-domain also produces abstractions that become meaningful again when translated back into the source-domain. This makes RL a particularly useful case for observing what is preserved, what is transformed, and what becomes newly legible through repeated crossings.

I then introduce several “supermemes” from translation studies, meaning recurring, widely shared assumptions about what translation is and how it ought to be done. They are used as an interpretive vocabulary for describing how ideas acquire new constraints when they are transferred between the source-domain of biology and the target-domain of computer science.

The goal of this writing is pragmatic: to assemble a shared vocabulary for thinking about how ideas travel between domains. By referencing translation theory, I advocate for this type of exchange to be read as a series of decisions that any translation must negotiate: what counts as faithful, where untranslatability is encountered or even becomes productive, and when a freer strategy is more effective than literal imitation. This vocabulary is intended to be usable across communities, offering artists, biologists, and computer scientists a common language to examine decisions to preserve, compress, omit for the sake of conceptual coherence and functioning in its new target context.

A CASE STUDY: REINFORCEMENT LEARNING AND DOPAMINE

A central move in twentieth-century learning research was to treat behavior as shaped by experience-dependent change under reinforcement, and to ask what internal variables such a process would require. In broad strokes, the psychological “source text” emphasized three recurring motifs. First, learning unfolds over time and is sensitive to temporal order. Second, outcomes modify future tendencies to act. Third, prediction matters, because learning is shaped not only by rewards but also by violations of expectation and other discrepancies between anticipated and observed events (Rescorla, 1988; Holland, 2016).

An explicit bridge between Pavlovian conditioning and instrumental action was articulated in two-process learning theory, which examined how Pavlovian associations can mediate or modulate instrumental behavior (Rescorla & Solomon, 1967).

When these motifs enter the target-domain of computer science, they do so through a particular representational choice. The learning problem is written in a grammar of *states*, *actions*, and *returns*. The target-domain then insists on additional commitments: an agent must be able to represent what is currently relevant (current *state*), choose what to do (typically following some learned

policy), and evaluate the long-term consequences of its actions (*returns* and *value*). This is already a translation, because a “state” is not an anatomical object and a “reward” is not itself a neurotransmitter. They are target-domain variables that render a broad class of learning problems mathematically explicit.

An important early step in this translation was to treat learning systems as adaptive networks that develop expectations and predictions from experience. Sutton and Barto framed this as a meeting point between psychological learning theory and adaptive system theory, emphasizing prediction as a central organizing concept (Sutton & Barto, 1981).

Reinforcement learning: a brief primer

RL is often introduced as an engineering framework, but its core objects can be read as a target-domain translation of questions that animal learning and behavioral psychology have asked for over a century: How does experience change future behavior? How do animals assign credit to actions when outcomes are delayed? How do expectations update when rewards arrive or fail to arrive? (Rescorla & Wagner, 1972; Sutton & Barto, 1981). Here, I want to provide a brief explainer on the central ideas in RL representations. This will allow a deeper understanding in our eventual arrival to the concept of Temporal-Difference (TD) learning, which turns out to be a guidepost to formulating a new hypothesis in dopamine signaling in the animal brain.

In an attempt to build something that can learn in a continuous environment (as opposed to an episodic one, an environment with a terminal state such as a game of Donkey Kong), RL formalizes learning from interaction between an agent and an environment. Time is typically modeled in discrete steps $t = 0, 1, 2, \dots$. At each step t , the agent observes a *state* S_t , selects an *action* A_t , receives a *reward* R_{t+1} , and transitions to a new state S_{t+1} .

The point is not that brains store a variable literally called “state.” Instead, S_t is a target-domain stand-in for the information available to the organism that matters for behavior: cues, contexts, and internal conditions that shape action selection in conditioning and instrumental learning paradigms. This is close in spirit to how learning theory treats “stimuli” and “contexts” as the relevant situation descriptors in experiments. In the RL formalism, the environment’s dynamics are summarized by transition probabilities:

$$\Pr(S_{t+1} = s' \mid S_t = s, A_t = a) = P(s' \mid s, a)$$

The Markov decision process (MDP) assumption says that the next situation depends on the current situation and action, not the full history. Biologically, this is a simplification. Its purpose is methodological. It makes it possible to define learning rules that can be analyzed and implemented. It corresponds to

the modeling move that “the relevant past can be compressed into a current context representation,” a move that is also implicit in many behavioral paradigms where trial history is treated as mediated through current cues and learned associations.

In behavioral psychology, the object of study is often the probability of a response given a stimulus or context. RL expresses that as a (commonly stochastic) policy:

$$\pi(a \mid s)$$

the probability of taking action a in situation s . This is a direct translation of “response tendencies” into a formal object that can be optimized. RL differs from supervised learning because there is no direct target label for each action, meaning that no action is labeled as optimal or suboptimal at first; instead, the agent must learn from delayed outcomes.

$$G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}, \text{ with } 0 \leq \gamma < 1,$$

This definition answers a target-domain question: what exactly is being maximized? The return formalizes the idea that behavior can be shaped by future consequences, not only immediate outcomes. The discount factor γ encodes how strongly future outcomes influence present choices. In the source-domain, sensitivity to delay and temporally displaced outcomes is widely studied under the umbrella of intertemporal choice and delay discounting, though the empirical forms and mechanisms can vary by species and task.

Value functions: expectations and learned predictions

A major family of learning theories treats learning as the acquisition and updating of expectations about relations among cues, actions, and outcomes, often expressed through prediction error as a driver of change. This emphasis is explicit in influential work on Pavlovian conditioning that argues conditioning is better understood as learning relations among events than as simple response transfer, and in review treatments that frame prediction error as a central organizing variable across modern associative learning theories (Pearce & Hall., 1980).

RL translates this stance into value functions, which make “expectation” into an explicit predictive quantity over return. The state-value function under a policy π is defined as:

$$V^{\pi}(s) = \mathbb{E}[G_t \mid S_t = s]$$

and the action-value function is:

$$Q^\pi(s, a) = \mathbb{E}_\pi[G_t \mid S_t = s, A_t = a]$$

These definitions do philosophical work as well as technical work. They assert that learning can be expressed in terms of internally meaningful predictive quantities, and they provide targets for algorithmic approximation in sequential decision problems.

Bellman equations: a consistency condition on prediction

Once value is defined as an expectation over future rewards, the Bellman expectation equation expresses a self-consistency property. If $V^\pi(s)$ is truly the expected discounted future reward from s , then it must equal expected immediate reward plus discounted expected value of the next state, under the same policy. One common form is:

$$V^\pi(s) = \sum_a \pi(a \mid s) \sum_{s'} P(s' \mid s, a) (\mathbb{E}[R_{t+1} \mid s, a, s'] + \gamma V^\pi(s'))$$

This equation provides the target-domain with a computational handle for learning and planning. It turns the vague goal “maximize reward over time” into a set of constraints that can be solved or approximated. Historically, this is exactly the point where the translation from the source-domain becomes transformative: the animal-learning intuition that “behavior is shaped by reinforcement and expectation” becomes an algorithmic object with clear update targets.

Crucially, these equations do not assert that brains literally solve Bellman equations. Instead, they express a structured hypothesis about what successful learning must approximate in any system that behaves as if it optimizes long-term consequences. The next section introduces temporal-difference learning as one influential algorithmic mechanism for approximating these consistency conditions using errors computed over time, and then follow how that target-domain mechanism becomes legible again as a hypothesis in neurobiology.

TEMPORAL-DIFFERENCE LEARNING

So far, the discussion has been comfortably reliant on a generous assumption: that informative reward signals arrive often enough, and cleanly enough, for value estimates to be anchored by directly observed outcomes. Many environments violate that assumption. In continuing tasks, and in many real-world settings that must be modeled as continuing, informative outcomes can be sparse, delayed, or buried in long stretches of ambiguity (Sutton, 1988).

Temporal-difference (TD) learning is the hinge where this translation becomes

especially instructive, because its motivation is not simply “biology-inspired.” It is an answer to concrete target-domain pressures about when learning signals are available and how credit should be assigned over time. Sutton’s 1988 paper introduced TD methods as incremental procedures specialized for multi-step prediction, and it emphasized a specific change in what counts as an error. Rather than computing error only by comparing a prediction against a final outcome, TD learning assigns credit using the difference between temporally successive predictions.

In its simplest form, TD learning defines a one-step temporal-difference error:

$$\delta_t = R_{t+1} + \gamma V(S_{t+1}) - V(S_t),$$

and updates the current value estimate by

$$V(S_t) \leftarrow V(S_t) + \alpha \delta_t,$$

with learning rate α .

This reorients the learning objective introduced in the previous section. Instead of treating learning as driven by alignment between $V(S_t)$ and the full return G_t , TD learning pushes $V(S_t)$ toward a bootstrapped target, $R_{t+1} + \gamma V(S_{t+1})$, so that each step supplies a locally available teaching signal. The algorithm is therefore organized around a predictive discrepancy that is defined at every transition, not only at the end of a trajectory. This character of TD learning, where an organism or agent can update its expectations using an error defined by the mismatch between what was expected and what is now implied by the next moment, helped make TD-style errors legible to later work in neuroscience.

Early computational neurobiology proposed that dopaminergic systems could distribute a teaching-like signal that resembles a prediction error over time (Montague, Dayan, & Sejnowski, 1996). Subsequent influential work argued that phasic dopamine activity in primates reflects errors in reward prediction and linked these dynamics explicitly to temporal-difference formulations (Schultz, 1992; Schultz, 1997; Hollerman & Schultz, 1998).

Empirical and modeling studies then tested and refined the quantitative fit between TD-style prediction errors and dopamine responses (Bayer, 2005). In humans, TD-derived prediction-error regressors were adopted as a standard interpretive tool in neuroimaging studies of reward learning (O’Doherty, 2003; Pessiglione, 2006), connecting target-domain learning rules back to neural signals in striatum and related circuits. Reviews in computational cognitive neuroscience subsequently consolidated and critiqued this cross-domain mapping, clarifying both what TD-style errors explain well and where the

analogy strains (Dayan & Niv, 2008; Glimcher, 2011).

Critiques of the TD-Learning Crossover

Long before RL's TD learning is linked to dopaminergic activity in computational neuroscience, researchers working in Biology had already raised objections to treating reinforcement as a universal explanatory currency for animal learning. Some questioned if reinforcement is even necessary for learning. Classic demonstrations of latent learning and "cognitive maps" argued that animals can acquire structured knowledge about environments without immediate reward, and that reward can modulate performance rather than create learning from scratch (Tolman, 1948). Other critiques emphasized that learning is not equally distributed for all stimulus-consequence pairings. For example, an experiment with rats tested the learning in rats by linking a "tasty" water that causes nausea to the neutral stimuli, bright light and loud noises. They found that the learned aversion attached far more strongly to the taste cue than to the audiovisual cue (Eelen, 2018). Work on selective association showed that organisms are biologically prepared to learn some relations far more readily than others, challenging the idea of a general-purpose reinforcement mechanism that is indifferent to content (Garcia & Koelling, 1966; Seligman, 1970; Domjan & Galef, 1983). Still others highlighted how species-typical action patterns can override reinforcement contingencies, revealing that what is learnable is constrained by evolved motor repertoires and ecological niches (Breland & Breland, 1961). In avoidance and threat learning, Bolles (1970) argued that what appears to be reinforcement-driven learning may often be better understood through species-specific defense reactions, again placing innate structure at the center rather than treating reward and punishment as sufficient sculpting tools. Finally, timing-focused accounts argued that temporal structure and informativeness are not peripheral parameters but core dimensions of what animals learn, creating pressure on models that treat time as bookkeeping rather than representation (Gallistel & Gibbon, 2000).

The critiques above are by no means a dismissal of RL as a discipline, but more as a variety of ideas that illustrates the fuzziness that inherently exists in the source-domain. Like all field of studies, there is infighting about everything from under the sun. These internal arguments and tensions surface when concepts are translated, carried across, into the target-domain. The re-writing of biological learning as value, policy, and error prediction formula created a new set of tools that could be sharpened. TD learning is a prime example of such a sharpening. Under the new constraints in the target-domain, programmers poses requirements such as incremental updates, long horizons, and explicit credit assignment in their system design: the need for learning signal to be defined at each transition rather than reserved for terminal outcomes gave rise to the change of learning objective from value prediction to error prediction, leading us directly to TD learning. Once formalized, TD learning returned to the source-domain in a new role. It became a quantitative interpretive framework or neural data, ost famously in

the proposal that phasic dopamine responses resemble a reward-prediction-error-like teaching signal (Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997). Even where the mapping is partial, contested, or later refined, the back-translation demonstrates a concrete benefit of cross-domain exchange: ideas that began as behavioral principles were recast as algorithms, and those algorithms then helped clarify the underlying logics of specific biological signals.

This is where translation theory becomes more than just a metaphor. The animal learning to RL algorithms to RPE signal hypothesis ley line is created by a series of choices about what is kept, what is compressed, and what liberties are taken for the sake of legibility in the target-domain. The next section turns to five central ideas from translation studies as a compact vocabulary for naming translational decisions involved in translating biological mechanism into the computational domain.

SUPERMEMES IN CROSS-DOMAIN TRANSLATION

Translation as a practice has been developing since ancient Mesopotamia, as evidenced by translation-adjacent artifacts such as the bilingual Sumerian–Akkadian lexical tablets. For much of human history, we have been juggling the challenges of carrying ideas across different domains of knowledge. Part of this practice involves defining the scope of what translation is, asking whether equivalence is possible or even desirable, and identifying the kinds of distortion that may be structurally unavoidable. Many of these discussions recur and evolve across periods and schools long before translation practices are formalized into theories.

To generalize these recurrences, translation theorist Andrew Chesterman proposed the term supermemes (Chesterman, 2016). A supermeme is built on the idea of a “meme,” meaning transmissible ideas: concepts that spread, stabilize, and mutate through a community. In Chesterman’s coinage, supermemes are not individual techniques or stylistic preferences, but high-level conceptual clusters that organize many smaller arguments in the “meme pool” of translation theory.

In what follows, I would like to introduce five supermemes from translation theory as a language framework for the act of translating concepts from one knowledge domain to another. The aim is not to claim that translating between languages and translating between disciplines are identical acts. Rather, by understanding each supermeme, it becomes possible to analyze an interpretive trajectory (for instance, the case study of animal learning to algorithms to computational neurobiology) as a sequence of translational decisions, rather than as loosely defined influences between disciplines.

Source-Target

The first supermeme is the most basic description of a translation: the assumption that a translation carries meaning from a source to a target. There

is almost a politics to this dichotomy.

In language translation, there is a tendency to interpret the source as the original and the target as derivative. But as we have seen in the RL case study, the roles in the source-target domain exchange can be fuzzy and circular. RL begins as a target-domain translation of animal learning and behavioral psychology: broad learning phenomena are rewritten into states, actions, rewards, and value. But once those target-domain objects exist, they become exportable back toward the source-domain as explanatory tools. TD-style errors and value-based learning provide a compact language for describing aspects of biological learning signals, even if only partially and under contestable assumptions. That is not to say that “dopamine computes δ_t in the TD-error algorithm”, saying that would risk collapsing the exchange into a category error. The source-target distinction helps in defining the relationship between TD-error and RPE a relational one rather than an identitarian one. A more careful framing is to say that “under a particular task formalization, the source-domain TD error predicts a pattern that resembles phasic dopamine responses in the target-domain.” A computational object that successfully fits the recorded data of a biological mechanism does not automatically mean that we should conclude that “biology is implementing the algorithm”, or that the computational object as the original. The computational object here offers a candidate description of what the biological system may be approximating, at some level and under some conditions.

Another interesting aspect in source-target dichotomy is the constraints imposed by the target-domain. When a biological source is translated into computation, the target-domain forces commitments about representation and function: what will count as a “state,” an “action,” a “reward,” and what will count as successful learning. These new constraints in formalism determines what the translation preserves, compresses, and discards. A computational model of the learning phenomenon often need to prioritize legibility, tractability, and operational definitions, even when these priorities do not align with biological realism. This contributes to the innate desire in scholars to defend the definitions of source-objects established in their own disciplines when it is being used in another discipline. For example, RL typically treats reward and value as scalar variables, whereas biological learning signals are often distributed across circuits and may encode multiple dimensions beyond a single common currency.

In practice, this supermeme supplies a diagnostic question that will recur throughout the remaining supermemes: what exactly is being treated as the source, and what is being demanded of the target? Once those are named, the next question becomes unavoidable and more specific: equivalence in what sense? That is the second supermeme.

Equivalence

This supermeme is the singular most anxiety-inducing concepts in the course

of translation: what counts as “sameness” across the gap. In language translation, translators have always negotiated “sameness” across languages. Before it was formalized and studied, though, it was a craft norm, a practical concern rather than a theoretical one. Some early perspectives in scholarly work introduced this overarching goal as “equivalent messages in two different codes” (Jakobson, 1959) and “the replacement of textual material in one language by equivalent textual material in another language” (Catford, 1965).

From that period forward, various schools in breaking down the idea of “sameness” became more prevalent, translation theorists were concerned with the type of “sameness” one can pursue. Equivalence could be pursued as similarity of linguistic form, or as similarity of effect/response in the receiving context in the binary of formal equivalence and dynamic equivalence (Nida, 1964). Later work focuses on defining equivalence with the target-cultural context in mind, overt vs. covert translation contrasts translations that openly retain the source text’s historical and cultural situatedness with those designed to function as if they were originally written for the target culture, often using a cultural filter to achieve pragmatic naturalness (House, 1981). Imitative vs. functional translation similarly separates resemblance-based fidelity, which prioritizes reproducing recognizable source patterns, from purpose-based fidelity, which prioritizes the equivalence in delivering the text’s intended purpose in the receiving context (Jakobsen, 1994a). That is all to illustrate that equivalence is rarely a simple claim that two strings of words match, it is more a multidimensional gradient that one must constantly balance when translating. We could dissect equivalence as a multitude in terms of function, effect, type, and reference; there is not a single “winning” definition.

In cross-discipline translation, the concept of equivalence becomes unstable since biology and computer science do not share the same representational substrate as translation between written languages. A biological mechanism can be described in multiple fidelity (molecular, cellular, behavioral, etc), whereas computational learning models are constructed by specifying an objective function and updating parameters via optimization procedures grounded in mathematical tools such as calculus and linear algebra. This makes it easy to smuggle in an unspoken equivalence claims. For example, saying “RL captures animal learning” can mean anything from the behavior looks similar to the same computations are implemented to the same objective is being optimized. The supermeme of equivalence is useful here because it forces that ambiguity into the open.

In analyzing the RL case, we can take a look at its equivalence through two equivalence subtypes, though you can most certainly analyze the case study using other subtypes that fits what your goal might be:

1. **Function equivalence:** the target-domain object plays a similar role in

explaining or predicting outcomes, even if the implementation differs.

2. **Mechanistic equivalence:** the target-domain object corresponds to a specific biological mechanism, including how information is represented and how computations are carried out in the substrate.

Returning to our case study, the equivalence between neuronal dopamine activity and TD-error is only possible when we reason through function equivalence whereas the mapping breaks down through the lens of mechanistic equivalence. The brain doesn't use the same representation model of states, discount factors, and value function that appear in the algorithm. The notion of equivalence holds up when we define it under the functional lens, where the sameness is defined with the outcome of learning from experience.

So different slicing of equivalence makes or breaks any translation. This brings us to a core point of conflict between the two disciplines: biologists and computer scientists often argue past one another because they are using different equivalence standards without naming them. A biologist may reject an RL model because it is not mechanistically faithful, while a computer scientist may defend it because it is functionally effective. Translation theory does not automatically adjudicate who is "right," but it makes the disagreement legible. Equivalence, specifically, reframes the question from "is the model correct?" to "correct in what sense, for what purpose, and at what level of description?"

Untranslatability

If equivalence names the qualities that can be carried across, then this supreme would be its rightful antithesis. No translation can preserve all dimensions from source to target. More often than not, when something is retained, another is sacrificed in the process. When you choose to preserve the phonetic rhythms of a poem in the target language, some of the wording and meaning might need to be reinvented, thus nudged further from the source. Untranslatability doesn't point toward a dead end, it is an opportunity to make peace with the unyielding nature of encoding and decoding in the natural language space.

In cross-domain translation, untranslatability shows up early because biology and computer science do not share a common ontology. The animal body's metabolic constraints, the biochemical interactions in neuromodulators, the heterogeneity of neural populations, these are all barriers to achieving mechanistic equivalence with TD-error. The TD-error, on the other hand, depends on a decomposition of experience into a discrete state S_t , a measurable reward R_{t+1} and an arbitrary discount factor γ . Some of the mismatch can be negotiated through different anatomization of equivalence. As translators, or model architects, we analyze and name the irreconcilability between the two domains, and decide what is a productive interpretation of

the agent's state given our learning objective, and how to update the agent's internal model of the world given the temporal horizon of the intended task.

Seen this way, many familiar critiques of RL from the biological community can be restated as claims of untranslatability. When critics argue that brains do not operate with a single common currency, or that learning signals multiplex multiple variables, they are not only disputing a model's accuracy. They are pointing to a mismatch in what the two domains can preserve under translation. Some of the mismatch can be addressed by expanding upon the target domain's formalism such as computing with the assumption of partially observable state, using deep neural networks to enrich heterogeneity of internal representation space, and increasing the number of learning objective, to name a few. But other mismatch might be structural. In RL's case, the target domain gains tractability by compressing the source-domain object.

Free vs. Literal

The free versus literal debate is possibly the fault lines that has caused the most amount of violence and harm in the Western tradition. Rooted in sacred text translation, the stakes of "accuracy" was not only linguistic, but also theological and juridical. One of the early figures in this debate is Eusebius Hieronymus, or commonly known as Saint Jerome, one of the pioneers that advocated for "sense for sense" translation over the "word for word" method. Yet even he carved out an exception for Sacred Scripture, claiming that "the very order of the words is a mystery" when it comes to a sacred object. This contradiction shows how literalism existed as more than a style. It was a posture of risk management. When a community treats a source text as untouchable, a word-for-word strategy can function as a shield against accusations of heresy.

The violence of that climate is not merely metaphorical. One frequently cited emblem is Étienne Dolet, whose translation choices were read as doctrinal claims. In a well-known episode surrounding his translation of Plato's *Axiochus*, adding a small phrase that implied the soul's non-immortality helped trigger prosecution (Pym, 1992); in later retellings, this "last three words" becomes the hinge that links a translation decision to trial and execution in 1546. Whether or not any single wording "caused" the outcome, the case remains instructive for the history of translation norms: literalism can be demanded not because it is the most illuminating way to carry meaning across, but because it is the safest way to avoid being punished for having carried meaning across.

This genealogy matters for bio-inspired computing because a similar moral atmosphere can sometimes gather around the source-domain of biology. A living system is imagined as intrinsically optimized, an authoritative standard, a sacred original not be with "tamper with". The literalist impulse demands that bio-inspired computer algorithms must stay faithful to its "original". Under this

kind of sensibility, deviations from biological realism is a failure of fidelity, even if the motivation between different systems might be explicitly different.

Translation as a discipline eventually valorized literary translators as creative agents who can freely negotiate the internal tensions in the source text in order to create a target text that actually serves the target culture. And as we saw in the RL case study, a free translation style offered scientists an opportunity to rewrite the learning mechanism so that the system can function with different constraints and for different purposes. Despite the imperfect equivalence, the target object evolved into a new objective that becomes useful to the source domain after back-translation.

The act of carrying biological mechanism into computer algorithms must omit, compress, formalize; it cannot carry everything across, nor should it be obligated to. In this framing, the real question is not whether the translation is literal enough, but whether its degree of freedom from the source is bounded by the purpose of the target context.

CONCLUSION

This essay has been an attempt to reframe bio-inspired computing as a translation rather than an imitation or a copy. The RL case study was introduced as a case of bidirectionality in cross-domain translation and a narrative device in advocating for free translation between disciplines. The supermemes offered a compact vocabulary to name the decisions that are otherwise left implicit.

Source-target clarifies each domain's role and responsibility. Equivalence strips absoluteness from the standard dictionary definition by introducing more nuanced comparisons. Untranslatability removes shortcomings as sites of compression and reinvention. Finally, free versus literal translation helps explain why some cross-domain arguments feel morally charged. When a source is treated as privileged and inviolable (whether this source is a sacred text or the sacred Nature), literalism can impose risks in translation. In bio-inspired modeling, however, responsible work often requires a freer strategy: one that is explicit about what is preserved, what is idealized, and why those departures are justified by the target context.

If there is a practical recommendation to draw from this theoretic framing, it is that bio-inspired algorithms should be evaluated with a clearer account of their translational stance. What is being translated from biology: a biochemical reaction, a behavioral feedback loop, or a representation model of a system? Which subtypes of equivalence are we trying to pursue: mechanistic, functional, stylistic, a hybrid? Where are the points of contention, the untranslatable bits? Clarifying these invisible choices can break us out of the dynamics of defaulting to defensiveness when faced with disciplinary disagreements, and instead, turn toward being inquisitive.

REFERENCES

- Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47(1), 129–141. doi:10.1016/j.neuron.2005.05.020
- Bolles, R. C. (1970). Species-specific defense reactions and avoidance learning. *Psychological Review*, 77(1), 32–48.
- Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, 16(11), 681–684.
- Catford, J. C. (1965). *A linguistic theory of translation: An essay in applied linguistics*. Oxford University Press.
- Chesterman, A. (2016). *Memes of translation: The spread of ideas in translation theory* (Rev. ed.). John Benjamins.
- Dayan, P., & Hinton, G. E. (1992). Feudal reinforcement learning. In *Advances in Neural Information Processing Systems* (pp. 271–278).
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18(2), 185–196.
- Domjan, M., & Galef, B. G., Jr. (1983). Biological constraints on instrumental and classical conditioning: Implications for general process theory. *American Psychologist*, 38, 410–421.
- Eelen, P. (2018). Classical conditioning: Classical yet modern. *Psychologica Belgica*, 58(1), 196–211. doi:10.5334/pb.451
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2), 289–344. doi:10.1037/0033-295X.107.2.289
- Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4(4), 123–124. doi:10.3758/BF03342209
- Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Suppl. 3), 15647–15654. doi:10.1073/pnas.1014269108
- Holland, P. C., & Schiffino, F. L. (2016). Mini-review: Prediction errors, attention and associative learning. *Neurobiology of Learning and Memory*, 131, 150–156.
- Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1(4), 304–309.
- House, J. (1981). *A model for translation quality assessment* (2nd rev. ed.). Gunter Narr Verlag.
- Jakobsen, A. L. (1994). Starting from the (other) end: Integrating translation and text

- production. In C. Dollerup & A. Lindegaard (Eds.), *Teaching translation and interpreting 2: Insights, aims, visions* (pp. 143–150). John Benjamins.
- Jakobson, R. (1959). On linguistic aspects of translation. In R. A. Brower (Ed.), *On translation* (pp. 232–239). Harvard University Press.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16(5), 1936–1947.
- Nida, E. A. (1964). *Toward a science of translating: With special reference to principles and procedures involved in Bible translating*. E. J. Brill.
- O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, 38(2), 329–337.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532–552.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442(7106), 1042–1045. doi:10.1038/nature05051
- Pym, A. (1992). *Translation and text transfer: An essay on the principles of intercultural communication*. Peter Lang.
- Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43(3), 151–160.
- Rescorla, R. A., & Solomon, R. L. (1967). Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review*, 74(3), 151–182.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.
- Schultz, W. (1992). Activity of dopamine neurons in the behaving primate. *Seminars in Neuroscience*, 4, 129–138.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599. doi:10.1126/science.275.5306.1593
- Seligman, M. E. P. (1970). On the generality of the laws of learning. *Psychological Review*, 77(5), 406–418.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 9–44. doi:10.1023/A:1022633531479

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88(2), 135–170.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208.

CHAN, WUN TING

Wun Ting Chan is a NYC-based computer scientist, translator, and artist from Hong Kong. Chan has a B.F.A. in Visual Communications from Northwest College of Art, a B.S. in Computer Science from Evergreen State College, and she is currently pursuing a M.S. in Computer Science at the Courant Institute of Mathematical Sciences, New York University. She is a member of the Agentic Learning AI Lab advised by Dr. Mengye Ren. Chan's artistic practice aims to disrupt the status quo by injecting neologisms as texts and as art objects into the cultural corpus. In her most recent work, she was commissioned by the Bioarts Ethical Advisory Komission (BEAK) for the English-Simplified Chinese translation of *Human Germline Gene Editing is Bioart: an Open Letter to Lulu and Nana* by Dr. Adam Zaresky, a bioart-centric, cross-disciplinary critique of human germline editing.

Website: <https://ting.directory>