A Measure-Theoretic Description of the Intrinsic View of Embodied Agents

Nihat Ay
Wolfgang Löhr

SFI WORKING PAPER: 2013-03-011

A Measure-Theoretic Description of the Intrinsic View of Embodied Agents

Nihat Ay[1,3] Wolfgang Löhr[1,2]

April 4, 2013

[1]Max Planck Institute for Mathematics in the Sciences, Inselstrasse 22, 04103 Leipzig, Germany
[2]Universität Duisburg-Essen, Thea-Leymann-Strasse 9, 45117 Essen, Germany
[3]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA

Abstract

We consider a general model of the sensori-motor loop of an agent interacting with the world. Here, we assume a particular causal structure, mechanistically described in terms of Markov kernels. In this generality, we define two σ-algebras of events in the world that describe two respective perspectives: (1) the perspective of an external observer, (2) the intrinsic perspective of the agent. Not all aspects of the world, seen from the external perspective, are accessible to the agent. This is expressed by the fact that the second σ-algebra is a subalgebra of the first one. We show that, under continuity and compactness assumptions, the global dynamics of the world can be simplified without changing the internal process. This simplification can serve as a minimal world model that the system must have in order to be consistent with the internal process.

Keywords: sensori-motor loop, embodied agent, intrinsic perspective, external observer, σ-algebra.

Contents

1 Introduction: the intrinsic view of embodied agents

In recent years, the field of embodied cognition provided evidence that the embodiment and the situatedness of agents play a key role in the development of intelligent behaviour [PB07]. Here, the intrinsic view of an agent, which is the basis for its control, differs from the (extrinsic) view of an external observer. The intrinsic view is also closely related to Uexküll's notion of an *Umwelt* which summarises all aspects of the world that have an effect on the agent and can be affected by the agent [Uex34].

A predominant observation within many case studies of the field of embodied intelligence is that quite simple control mechanisms can exploit the constrains of the embodiment and thereby lead to very complex behaviour, seen from outside. This exploitation constitutes an important part of intelligent behaviour. Therefore, a formal understanding of these two perspectives is required for understanding intelligence.

In this paper, we consider a quite natural general model of the sensori-motor loop (see e.g. [TP, ZAD]) based on the theory of causal networks [Pea]. This model will allow us to formalise what we mean by intrinsic and extrinsic in terms of σ-algebras. These are basic mathematical objects from measure theory that describe a natural set of observables which are assigned to an observer. Already at this very general level, a description of the gap between the intrinsic and extrinsic perspective is possible. Clearly, this gap is not visible by the agent, and we will show that a corresponding intrinsic model of the world can be constructed by the agent as replacement of the actual mechanisms in the world without any structure in addition to the intrinsic structure.

2 A formal model of the sensori-motor loop

We assume basic knowledge from measure and probability theory and refer to the comprehensive volumes [Bog07a, Bog07b] on measure theory and to the textbook [Bau68] on probability theory.

In this section we provide a general model for the *sensori-motor loop* of an agent (C) interacting with the world (W). We assume that this interaction is mediated through the agent's sensors (S) and actuators (A). More precisely, in each instant of time the agent takes a measurement from the world through its sensors and affects the world through its actuators. Formally, we have random variables W_n, S_n, C_n, A_n, $n \in \mathbb{N}$, taking values in (measurable) spaces $\mathsf{W}, \mathsf{S}, \mathsf{C}, \mathsf{A}$. For technical reasons, we assume that these are Souslin spaces, equipped with their respective Borel σ-algebras $\mathfrak{B}(\mathsf{W})$, $\mathfrak{B}(\mathsf{S})$, $\mathfrak{B}(\mathsf{C})$, and $\mathfrak{B}(\mathsf{A})$. From our above motivation it is natural to assume that the variables have a joint distribution that factorizes according to the graph shown in Figure 1, where α, β, φ, and π, defined as Markov kernels, model the mechanisms of the sensori-motor loop.[1]

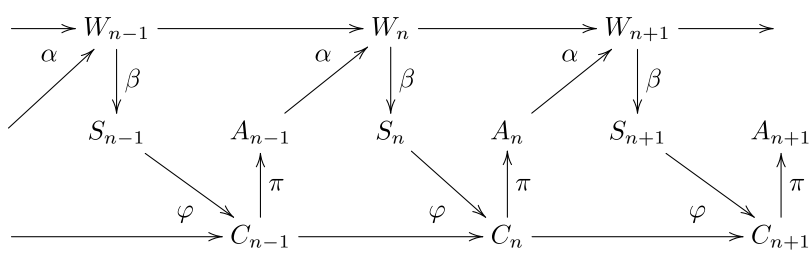

Figure 1

Together with an initial distribution of W_0, S_0, C_0, A_0, the Markov kernels α, β, φ, and π specify the distribution of the process. As shown in Figure 1, we denote the "world update kernel" from $\mathsf{A} \times \mathsf{W}$ to W by α and the "sensor kernel" from W to S by β. Let α_a, $a \in \mathsf{A}$, be the "world update" when the action is a, i.e. $\alpha_a(w) = \alpha(a, w)$. Later, we also use the "agent update kernel" φ from $\mathsf{C} \times \mathsf{S}$ to C and the "policy kernel" π from C to A.

In what follows, we address the following two natural problems:

1. Which structure in the world is used by the mechanisms of the sensori-motor loop?

2. Which structure of the world is visible from the intrinsic perspective of the agent?

We will show that these problems can be appropriately addressed by defining corresponding σ-algebras.

[1]The relation between factorization and conditional independence properties is subject of graphical models theory. See the standard reference [Lau96].

3 Minimal σ-algebras of the world

3.1 Minimal separately measurable σ-algebra

In order to address the above problems, we fix the Borel σ-algebras on the "agent part"—S, C, and A—of the system. Based on these internal σ-algebras, we consider various sub-σ-algebras of the Borel σ-algebra on W that describe agent related events in the world. In measure-theoretic terms, we study minimal σ-algebras on W that satisfy natural measurability conditions. The most natural ansatz is given by the distinctions that are possible only through sensor measurements. They correspond to the σ-algebra generated by the kernel β, that is $\sigma(\beta)$. However, this is not necessarily consistent with the world dynamics given by the Markov kernel α. Therefore, we consider the following measurability condition. We call a σ-algebra $\mathcal{W} \subseteq \mathfrak{B}(\mathsf{W})$ **(jointly) measurable** if both β and α remain measurable when W is equipped with $\mathcal{W}$ instead of $\mathfrak{B}(\mathsf{W})$. By general assumption, $\mathfrak{B}(\mathsf{W})$ is jointly measurable. It turns out that joint measurability is a quite strong condition. Therefore, it is natural to consider the following weaker measurability condition. We call $\mathcal{W}$ **separately measurable** if, when we equip W with $\mathcal{W}$, β is measurable and α is separately measurable in the sense that for every $a \in \mathsf{A}$, the Markov kernel

$$\alpha_a : (\mathsf{W}, \mathcal{W}) \to \mathcal{P}(\mathsf{W}, \mathcal{W}), \qquad w \mapsto \alpha(a, w),$$

is measurable. Note that, because α is Borel measurable, the functions $a \mapsto \alpha(a, w)$ are measurable for any $\mathcal{W} \subseteq \mathfrak{B}(\mathsf{W})$.

It is straight-forward to construct the unique minimal (w.r.t. partial ordering by inclusion) separately measurable sub-σ-algebra $\mathcal{W}_{\text{ext}}$ of $\mathfrak{B}(\mathsf{W})$.

Lemma 1. *Let $\mathcal{W}_0 := \sigma(\beta)$ and for $n \in \mathbb{N}$ define $\mathcal{W}_n$ recursively by*

$$\mathcal{W}_n := \sigma\big(\alpha_a \colon \mathsf{W} \to \mathcal{P}(\mathsf{W}, \mathcal{W}_{n-1}),\ a \in \mathsf{A}\big) = \sigma\big(\alpha_a(\,\cdot\,; B),\ \ a \in \mathsf{A},\ B \in \mathcal{W}_{n-1}\big)$$

Then $\mathcal{W}_{\text{ext}} := \sigma\big(\bigcup_{n\in\mathbb{N}} \mathcal{W}_n\big)$ is the unique minimal separately measurable σ-algebra, i.e.

$$\mathcal{W}_{\text{ext}} = \bigcap \big\{ \mathcal{W} \subseteq \mathfrak{B}(\mathsf{W}) \;\big|\; \mathcal{W} \ \sigma\text{-algebra, } \beta \text{ measurable, } \alpha_a\ \mathcal{W}\text{-}\mathcal{W}\text{-measurable for all } a \in \mathsf{A} \big\}$$

Proof. "$\subseteq$": Clearly, any σ-algebra $\mathcal{W}$ from the set on the right-hand side has to contain $\mathcal{W}_0$ because β is $\mathcal{W}$-measurable. Further, if it contains $\mathcal{W}_{n-1}$, it also has to contain $\mathcal{W}_n$, because α_a must be measurable for all $a \in \mathsf{A}$. Thus it contains $\mathcal{W}_{\text{ext}}$.

"$\supseteq$": We have to show that $\mathcal{W}_{\text{ext}}$ is separately measurable. β is measurable, because $\mathcal{W}_{\text{ext}} \supseteq \mathcal{W}_0 = \sigma(\beta)$. $\bigcup_n \mathcal{W}_n$ is an intersection stable generator of $\mathcal{W}_{\text{ext}}$, thus for measurability of α_a, it is sufficient that $\alpha_a(\,\cdot\,; B)$ is $\mathcal{W}_{\text{ext}}$-measurable for $B \in \bigcup_n \mathcal{W}_n$. But by definition of $\mathcal{W}_n$, $\alpha_a(\,\cdot\,; B)$ is $\mathcal{W}_n$-measurable for $B \in \mathcal{W}_{n-1}$. □

Note that there is no reason why α should be jointly measurable when we equip W with $\mathcal{W}_{\text{ext}}$. When we are working with a separately but not jointly measurable σ-algebra, we are rather working with a family $(\alpha_a)_{a\in\mathsf{A}}$ of kernels than with a single kernel α. We do not know if a minimal jointly measurable σ-algebra exists in general. The above construction does not work well for α instead of α_a, because we want a product σ-algebra on $\mathsf{A} \times \mathsf{W}$ and taking products is not compatible with intersections in the sense that $\mathcal{A} \otimes (\mathcal{W} \cap \mathcal{W}') \subsetneq (\mathcal{A} \otimes \mathcal{W}) \cap (\mathcal{A} \otimes \mathcal{W}')$ in general. Of course, every jointly measurable σ-algebra is separately measurable and thus has to contain $\mathcal{W}_{\text{ext}}$. Also note that $\mathcal{W}_{\text{ext}}$ need not be countably generated, which might cause technical problems when working with $\mathcal{W}_{\text{ext}}$. Next we show that in the "nice case" where $\mathcal{W}_{\text{ext}}$ is countably generated, α is jointly measurable.

Proposition 2. *If $\mathcal{W}_{\text{ext}}$ is countably generated, then it is jointly measurable, and in particular the unique minimal jointly measurable σ-algebra.*

Proof. Let $\mathcal{A} := \mathfrak{B}(\mathsf{A})$. We have to show that $f\colon \mathsf{A} \times \mathsf{W} \to [0,1]$, $(a,w) \mapsto \alpha(a,w;B)$ is $(\mathcal{A} \otimes \mathcal{W}_{\text{ext}})$-measurable for arbitrary choice of $B \in \mathcal{W}_{\text{ext}}$. Because $\mathcal{W}_{\text{ext}}$ is countably generated, $\mathcal{A} \otimes \mathcal{W}_{\text{ext}}$ is a countably generated sub-σ-algebra of the Borel σ-algebra of the Souslin space $\mathsf{A} \times \mathsf{W}$. It follows from Blackwell's theorem (see Appendix A) and the fact that f is Borel measurable that f is $\mathcal{A} \otimes \mathcal{W}_{\text{ext}}$ measurable if and only if it is constant on the atoms of $\mathcal{A} \otimes \mathcal{W}_{\text{ext}}$. The atoms are obviously of the form $\{\,a\,\} \times F$, where $a \in \mathsf{A}$ and F is an atom of $\mathcal{W}_{\text{ext}}$. Because α_a is measurable w.r.t. $\mathcal{W}_{\text{ext}}$, $f(a,\cdot\,)$ is constant on the atom F. Thus, f is constant on $\{\,a\,\} \times F$ and therefore jointly measurable. ☐

A simple sufficient condition for $\mathcal{W}_{\text{ext}}$ to be countably generated is that there are only countably many possible actions, i.e. A is countable.

Corollary 3. *Let A be countable. Then $\mathcal{W}_{\text{ext}}$ is countably generated and jointly measurable.*

Proof. Because $\mathfrak{B}(\mathsf{S})$ is countably generated, $\mathcal{W}_0$ is countably generated. If $\mathcal{W}_{n-1}$ is countably generated, the same holds for $\sigma\big(\alpha_a\colon \mathsf{W} \to \mathcal{P}(\mathsf{W}, \mathcal{W}_{n-1})\big)$ for any $a \in \mathsf{A}$. Because A is countable, $\mathcal{W}_n$ is generated by a countable union of countably generated σ-algebras, thus it is countably generated, and the same holds for $\mathcal{W}_{\text{ext}}$. ☐

3.2 A countably generated, almost jointly measurable σ-algebra in the memoryless case

In this section, *we assume that the agent is memoryless, i.e. C_n is conditionally independent of C_{n-1} given S_n*. Then we can concatenate the kernels from W to S, from S to C, and from C to A to obtain a new kernel γ from W to A. We then have the following situation, where the C and S components are marginalised (integrated) out.

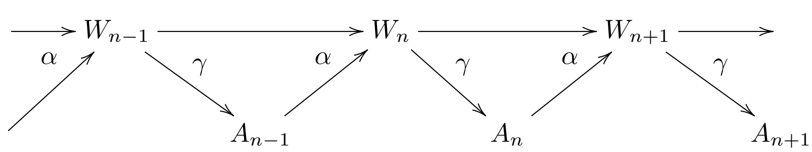

Note that if β is $\mathcal{W}$-measurable, the same holds for γ, but the converse need not be true. We introduce yet another kernel κ which is the combination of γ and α, i.e. $\kappa\colon \mathsf{W} \to \mathcal{P}(\mathsf{A} \times \mathsf{W})$, $\kappa(w) = \gamma(w) \otimes \alpha(\,\cdot\,, w)$. The reason to do this is that while kernels mapping *from* a product space complicate finding minimal σ-algebras (with product structure), this is not the case for kernels mapping *into* a product space. The variables W_n, A_n, $n \in \mathbb{N}$, factorise also according to the following graphical model.

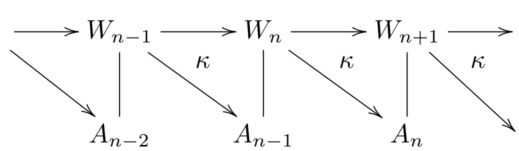

We can define the minimal σ-algebra $\mathcal{W}_\kappa$ such that β and κ are measurable in the same way as we defined $\mathcal{W}_{\text{ext}}$.

Lemma 4. *Let $\mathcal{W}'_0 := \sigma(\beta)$ and*

$$\mathcal{W}'_n := \sigma\big(\kappa\colon \mathsf{W} \to \mathcal{P}(\mathsf{A} \times \mathsf{W}, \mathcal{A} \otimes \mathcal{W}'_{n-1})\big) = \sigma\big(\kappa(\,\cdot\,; B),\ B \in \mathcal{A} \otimes \mathcal{W}'_{n-1}\big).$$

Then $\mathcal{W}_\kappa := \sigma\big(\bigcup_{n\in\mathbb{N}} \mathcal{W}'_n\big)$ is the unique minimal σ-algebra on W s.t. β and κ are measurable. Furthermore, $\mathcal{W}_\kappa$ is countably generated.

Proof. Analogous to the proof of Lemma 1 and Corollary 3. ☐

In the following, consider κ as kernel from $(\mathsf{W}, \mathcal{W}_\kappa)$ to $(\mathsf{A} \times \mathsf{W}, \mathcal{A} \otimes \mathcal{W}_\kappa)$. Note that because $\mathcal{W}_\kappa$ is countably generated, the quotient space obtained by identifying atoms of $\mathcal{W}_\kappa$ to points is again a Souslin space[2] (see Appendix A). Thus it is technically nice and, in particular, we can factorise κ into γ and some kernel α' from $\mathsf{A} \times \mathsf{W}$ to W by choosing regular versions of conditional probability. Then α' is jointly measurable. The draw-back is that $\alpha'(\,\cdot\,, w)$ is only defined $\gamma(w)$-almost surely and we cannot guarantee that $\alpha' = \alpha$ is a valid choice, i.e. that α is $\mathcal{A} \otimes \mathcal{W}_\kappa$-measurable. We easily get the following.

Lemma 5. *$\mathcal{W}_\kappa$ is the unique minimal σ-algebra on W that satisfies the following condition. β is measurable and there exists a (jointly measurable) kernel α' from $\mathsf{A} \times \mathsf{W}$ to W, s.t., for every measure $\mu \in \mathcal{P}\big(\mathsf{W}, \mathfrak{B}(\mathsf{W})\big)$, $\alpha = \alpha'$ $(\mu \otimes \gamma)$-almost surely.*

Proof. The above discussion shows that $\mathcal{W}_\kappa$ satisfies the condition (note that we can w.l.o.g. assume that μ is a Dirac measure). If, on the other hand, α' as above exists, then κ equals the composition of γ and α'. In particular, κ is measurable and Lemma 4 yields the claim. □

The condition $\alpha = \alpha'$ a.s. w.r.t. every measure of the form $\mu \otimes \gamma$ means that the difference between α and α' is not visible regardless of any changes we might impose on the environment. The situation, however, may change if the agent changes its action policy π, thereby changing the kernel γ. Then the difference between α and α' can become important and α' as well as $\mathcal{W}_\kappa$ would have to be changed.

We trivially have that every jointly measurable σ-algebra $\mathcal{W}$ on W must contain $\mathcal{W}_\kappa$. In particular, if $\mathcal{W}_{\text{ext}}$ is countably generated, $\mathcal{W}_\kappa \subseteq \mathcal{W}_{\text{ext}}$. This is probably not true in general.

4 The world from an intrinsic perspective

4.1 Sensory equivalence

In what follows we use equivalence relations to coarse grain the world states and apply the constructions described in Appendix A.

Denote by $P_S^{a_\mathbb{N}}(w) \in \mathcal{P}(\mathsf{S}^\mathbb{N})$ the distribution of sensor values when the "world" is initially in the state $W_1 = w \in \mathsf{W}$ and the agent performs the action sequence $a_\mathbb{N} \in \mathsf{A}^\mathbb{N}$, i.e. $A_n = a_n$. That is, we modify the agent policy π in a time-dependent way such that π is replaced by $(\pi_n)_{n\in\mathbb{N}}$ and π_n ignores the memory (and thus the sensors) and outputs (the Dirac measure in) the value a_n. The sensor and world-update kernels β and α, however, remain unchanged. More explicitly, for $B = B_1 \times \cdots \times B_n \times \mathsf{S} \times \cdots \in \mathfrak{B}(\mathsf{S}^\mathbb{N})$,

$$P_S^{a_\mathbb{N}}(w_1)(B) = \int \cdots \int \beta(w_1;\, B_1) \cdots \beta(w_n;\, B_n)\ \alpha_{a_{n-1}}(w_{n-1};\, \mathrm{d}w_n)\ \cdots\ \alpha_{a_1}(w_1;\, \mathrm{d}w_2).$$

Now we define an equivalence relation $\sim_{\mathrm{s}}$, called *sensory equivalence*, on W by

$$w \sim_{\mathrm{s}} w' \qquad :\Leftrightarrow \qquad P_S^{a_\mathbb{N}}(w) = P_S^{a_\mathbb{N}}(w')\ \ \forall a_\mathbb{N} \in \mathsf{A}^\mathbb{N}.$$

More generally, we obtain the *intrinsic σ-algebra*

$$\mathcal{W}_{\text{int}} := \sigma\Big(P_S^{a_\mathbb{N}},\ a_\mathbb{N} \in \mathsf{A}^\mathbb{N}\Big),$$

which describes the information about the world that can in principle be obtained by the agent through its sensors. Obviously, the atoms $[\cdot]_{\mathcal{W}_{\text{int}}}$ of the intrinsic σ-algebra are given precisely by the sensory equivalence, i.e. $[w]_{\mathcal{W}_{\text{int}}} = \{\, w' \in \mathsf{W} \mid w' \sim_{\mathrm{s}} w \,\}$ for all $w \in \mathsf{W}$.

Lemma 6. *$\mathcal{W}_{\text{int}} \subseteq \mathcal{W}_{\text{ext}}$. In particular*

$$[w]_{\mathcal{W}_{\text{ext}}} \subseteq [w]_{\mathcal{W}_{\text{int}}} = \{\, w' \in \mathsf{W} \mid w' \sim_{\mathrm{s}} w \,\} \qquad \forall w \in \mathsf{W}.$$

[2]more precisely there exists a Souslin topology such that the Borel σ-algebra coincides with the final σ-algebra induced by the canonical projection from $(\mathsf{W}, \mathcal{W}_\kappa)$ onto the quotient

Proof. β and α_a are measurable w.r.t. $\mathcal{W}_{\text{ext}}$. Because the σ-algebra on $\mathcal{P}(\mathsf{S}^{\mathbb{N}})$ is generated by the evaluations, and the cylinder sets form a generator of the σ-algebra on $\mathsf{S}^{\mathbb{N}}$, measurability of β and α_a implies measurability of the function $P_S^{a_{\mathbb{N}}}$ for every $a_{\mathbb{N}} \in \mathsf{A}^{\mathbb{N}}$. Hence $\mathcal{W}_{\text{int}} \subseteq \mathcal{W}_{\text{ext}}$. This directly implies the corresponding inclusion for the atoms. ☐

Equality in the above lemma does not hold, as the following example shows.

Example 7. Let $\mathsf{W} := \{1, \ldots, 5\}$, $\mathsf{S} = \{0, 1\}$ and $|\mathsf{A}| = 1$, i.e. the agent is only observing (a state-emitting HMM). Let $\beta(1) = \beta(4) = \beta(5) = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1$, $\beta(2) = \delta_0$, and $\beta(3) = \delta_1$. Further let $\alpha(1) = \frac{1}{2}\delta_2 + \frac{1}{2}\delta_3$, $\alpha(2) = \alpha(3) = \alpha(4) = \delta_4$, $\alpha(5) = \delta_5$. α can be illustrated as

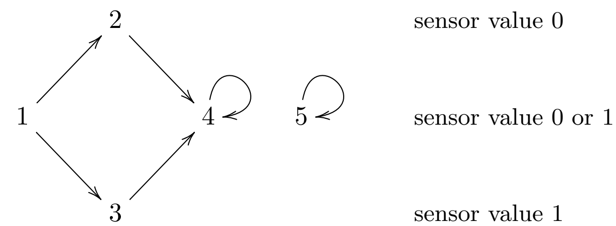

Then $1 \sim_{\text{s}} 4 \sim_{\text{s}} 5 \nsim_{\text{s}} 2, 3$ and $\text{At}(\mathcal{W}_0) = \big\{ \{1,4,5\}, \{2\}, \{3\} \big\}$, but

$$\text{At}(\mathcal{W}_{\text{ext}}) \;=\; \text{At}(\mathcal{W}_1) \;=\; \big\{ \{1\}, \{2\}, \{3\}, \{4,5\} \big\}.$$

Thus 1 and 4 are identified by $\sim_{\text{s}}$ because they produce identical sequences of sensor values, but they are not identified by $\mathcal{W}_{\text{ext}}$ because they have non-identified successors. The definition of $\mathcal{W}_{\text{ext}}$ requires that α remains unchanged, while the same sensor values can be produced with the partition given by $\sim_{\text{s}}$ by modifying α to α', where $\alpha'(1) = \alpha'(4)$ is an arbitrary convex combination of $\alpha(1)$ and $\alpha(4)$. ♢

4.2 Sensor-preserving modification of the world

Example 7 suggests that one might be able to interpret the coarser partition given by sensory equivalence as description of the relevant part of the world, provided one is allowed to modify the world update kernel α in such a way that the distribution of sensor values is preserved. Intuitively, one just has to choose one of the values α_a takes on a given $\sim_{\text{s}}$-equivalence class.

Of course these selections have to be done in a measurable way, and we need technical restrictions to deal with this problem. Namely, we assume that the world W is *compact*, and the sensor kernel β as well as the world update kernels $\alpha_a \colon \mathsf{W} \to \mathcal{P}(\mathsf{W})$ for every given action $a \in \mathsf{A}$ are *continuous*. As usual, $\mathcal{P}(\mathsf{W})$ is equipped with the weak topology induced by bounded continuous functions. Note that compactness and metrisability of W also implies compactness and metrisability of $\mathcal{P}(\mathsf{W})$.

Under these assumptions we can prove that it is possible to modify the world update (and with it the smallest separately measurable σ-algebra $\mathcal{W}_{\text{ext}}$) in such a way that the sensor process is preserved and equality holds in Lemma 6. Furthermore, the "new $\mathcal{W}_{\text{ext}}$" is countably generated and jointly measurable for the modified system.

Definition 8. Let $\alpha' \colon \mathsf{A} \times \mathsf{W} \to \mathcal{P}(\mathsf{W})$ be an "alternative" world update kernel.

1. We call α' **equivalent** to α if for every $w \in \mathsf{W}$ and $a_{\mathbb{N}} \in \mathsf{A}^{\mathbb{N}}$ the sensor process $P_S^{a_{\mathbb{N}}}(w)$ coincides with the sensor process $P_{S,\alpha'}^{a_{\mathbb{N}}}(w)$ obtained by replacing α with α'.

2. Denote by $\mathcal{W}_{\text{ext}}^{\alpha'}$ the smallest separately measurable σ-algebra of the system where α is replaced by α'.

Lemma 9. *Assume that α_a is continuous for every $a \in \mathsf{A}$ and β is continuous. Then $P_S^{a_{\mathbb{N}}}$ is continuous for every $a_{\mathbb{N}} \in \mathsf{A}^{\mathbb{N}}$.*

Proof. Easy to see directly or a special case of [Kar75, Thm. 1]. □

Proposition 10. *Let* W *be compact,* β *and* α_a *continuous for every* $a \in \mathsf{A}$*. Then there is a kernel* $\alpha'\colon \mathsf{A} \times \mathsf{W} \to \mathcal{P}(\mathsf{W})$*, such that* α' *is equivalent to* α *and*

$$\mathcal{W}_{\mathrm{ext}}^{\alpha'} = \mathcal{W}_{\mathrm{int}}. \tag{1}$$

In particular,

$$[w]_{\mathcal{W}_{\mathrm{ext}}^{\alpha'}} = \{\, w' \in \mathsf{W} \mid w' \sim_{\mathrm{s}} w \,\} \qquad \forall\, w \in \mathsf{W}.$$

Furthermore, $\mathcal{W}_{\mathrm{ext}}^{\alpha'}$ *is countably generated as well as jointly measurable (for the modified system with* α *replaced by* α'*).*

Proof.

1. Let $X := \mathcal{P}(\mathsf{S}^{\mathbb{N}})^{\mathsf{A}^{\mathbb{N}}}$ be the set of mappings from action sequences to distributions of sensor sequences, equipped with product topology. Given an initial state w of the world, denote by $F(w)$ the corresponding kernel from action sequences to sensor sequences, i.e. $F\colon \mathsf{W} \to X$, $F(w) = \big(a_{\mathbb{N}} \mapsto P_S^{a_{\mathbb{N}}}(w)\big)$. Note that F generates $\mathcal{W}_{\mathrm{int}}$, i.e. $\sigma(F) = \mathcal{W}_{\mathrm{int}}$. Because every $P_S^{a_{\mathbb{N}}}$ is continuous (Lemma 9) and X carries the product topology, F is a continuous function from the compact metrisable space W into the Hausdorff space X. In particular, the image $F(\mathsf{W})$ is also compact and metrisable.
2. We can apply a classical selection theorem, e.g. Theorem 6.9.7 in [Bog07b], and obtain a measurable right-inverse $G\colon F(\mathsf{W}) \to \mathsf{W}$ with $F \circ G = \mathrm{id}_{F(\mathsf{W})}$. Define $\varsigma := G \circ F$. Then ς is measurable, and $\sigma(\varsigma) \subseteq \sigma(F)$. On the other hand, $\sigma(F) = \sigma(F \circ G \circ F) \subseteq \sigma(\varsigma)$. Hence,
$$\sigma(\varsigma) = \sigma(F) = \mathcal{W}_{\mathrm{int}}. \tag{2}$$
Define $\alpha'_a := \alpha_a \circ \varsigma$ for every $a \in \mathsf{A}$.
3. A simple induction shows that α' is indeed equivalent to α: For $B = B_1 \times \cdots \times B_n \times \mathsf{S} \times \cdots \in \mathfrak{B}(\mathsf{S}^{\mathbb{N}})$ and $C := B_2 \times \cdots \times B_n \times \mathsf{S} \times \cdots$, we obtain by induction over n
$$P_{S,\alpha'}^{a_{\mathbb{N}}}(w)(B) = \int \beta(w; B_1) P_{S,\alpha'}^{a_{\{2,3,\dots\}}}(\cdot)(C)\, \mathrm{d}\alpha'_{a_1}(w) = P_S^{a_{\mathbb{N}}}\big(\varsigma(w)\big)(B) = P_S^{a_{\mathbb{N}}}(w)(B)$$
4. We claim that $\mathcal{W}_{\mathrm{ext}}^{\alpha'} = \sigma(\varsigma)$. Indeed, α'_a is $\sigma(\varsigma)$-measurable by definition, and β is $\sigma(\varsigma)$-measurable, because $\beta(w)$ is a marginal of $F(w)(a_{\mathbb{N}})$ for any $a_{\mathbb{N}}$. Therefore, $\sigma(\varsigma)$ is separately measurable in the modified system and $\mathcal{W}_{\mathrm{ext}}^{\alpha'} \subseteq \sigma(\varsigma)$. On the other hand, $P_S^{a_{\mathbb{N}}} = P_{S,\alpha'}^{a_{\mathbb{N}}}$ is $\mathcal{W}_{\mathrm{ext}}^{\alpha'}$-measurable, thus the same holds for F and $\sigma(\varsigma) = \sigma(F) \subseteq \mathcal{W}_{\mathrm{ext}}^{\alpha'}$. Hence $\mathcal{W}_{\mathrm{int}} = \mathcal{W}_{\mathrm{ext}}^{\alpha'}$ follows from (2).
5. Since ς is a function into a space with countably generated σ-algebra, $\mathcal{W}_{\mathrm{ext}}^{\alpha'} = \sigma(\varsigma)$ is countably generated. In particular, it is jointly measurable by Proposition 2. □

Remark (Non-compact W)**.** For a non-compact world W, we can still obtain an equivalent α' satisfying (1) if we relax the condition that α' needs to be Borel measurable. Instead, it is only universally measurable, i.e. μ-measurable for every $\mu \in \mathcal{P}(\mathsf{W})$. To see this, just replace the selection theorem used in the proof of Proposition 10 by a selection theorem for Souslin spaces, e.g. Theorem 6.9.1 in [Bog07b]. The drawback is that the universal σ-algebra is not countably generated and we do not obtain joint measurability of $\mathcal{W}_{\mathrm{ext}}^{\alpha'}$.

A Appendix: state reduction and quotient construction

Let $(X, \mathcal{F})$ be a measurable space. The **atom** of $\mathcal{F}$ containing $x \in X$ and the set of atoms of $\mathcal{F}$ are defined as

$$[x]_{\mathcal{F}} := \bigcap_{x \in F \in \mathcal{F}} F \qquad \text{and} \qquad \mathrm{At}(\mathcal{F}) := \big\{\, [x]_{\mathcal{F}} \;\big|\; x \in X \,\big\}.$$

Note that if $\mathcal{F}$ is countably generated, $[x]_{\mathcal{F}}$ is a measurable set, $[x]_{\mathcal{F}} \in \mathcal{F}$. In general, however, $[x]_{\mathcal{F}}$ need not be measurable. We recall Blackwell's theorem.

Blackwell's theorem. *Let X be a Souslin space and $\mathcal{F} \subseteq \mathfrak{B}(X)$ a countably generated sub-σ-algebra of the Borel σ-algebra. Then*

$$\mathcal{F} = \left\{ F \in \mathfrak{B}(X) \,\middle|\, F = \bigcup_{x \in F} [x]_{\mathcal{F}} \right\}.$$

Corollary 11. *Let X be a Souslin space, $\mathcal{F} \subseteq \mathfrak{B}(X)$ a countably generated σ-algebra, and $f \colon X \to \mathbb{R}$ measurable. Then f is $\mathcal{F}$-measurable if and only if it is constant on the atoms of $\mathcal{F}$.*

According to Blackwell's theorem, a countably generated sub-σ-algebra of a Souslin space X is uniquely determined by the set of it atoms. $\mathrm{At}(\mathcal{F})$ is a partition of X into $\mathfrak{B}(X)$-measurable sets. Note, however, that not every partition of X into measurable sets is the set of atoms of a countably generated sub-σ-algebra of $\mathfrak{B}(X)$.

Given any measurable space $(X, \mathcal{F})$, we can define the quotient space $X_{\mathcal{F}}$ as the set $\mathrm{At}(\mathcal{F})$ of atoms of $\mathcal{F}$ equipped with the final σ-algebra $\mathcal{X}_{\mathcal{F}}$ of the canonical projection $[\,\cdot\,]_{\mathcal{F}} \colon X \to \mathrm{At}(\mathcal{F})$. Then a set $B \subseteq X_{\mathcal{F}}$ of atoms is by definition measurable iff $\bigcup B = \bigcup_{[x]_{\mathcal{F}} \in B} [x]_{\mathcal{F}} \in \mathcal{F}$. Note that, obviously, $B \mapsto \bigcup B$ is a complete isomorphism of boolean algebras from $\mathcal{X}_{\mathcal{F}}$ onto $\mathcal{F}$. The following lemma follows easily from the standard theory of analytic measurable spaces and is one of the reasons why Souslin spaces, rather than Polish spaces, are the "right" class of spaces to work with in our setting.

Lemma 12. *Let X be a Souslin space and $\mathcal{F} \subseteq \mathfrak{B}(X)$ a sub-σ-algebra. Then $\mathcal{X}_{\mathcal{F}}$ is the Borel σ-algebra of some Souslin topology on $X_{\mathcal{F}}$ if and only if $\mathcal{F}$ is countably generated.*

Acknowledgement

Nihat Ay is grateful for stimulating discussions with Keyan Zahedi.

References

[Bau68] H. Bauer. *Wahrscheinlichkeitstheorie und Grundzüge der Maßtheorie.* Walter de Gruyter, 1968.

[Bog07a] V. I. Bogachev. *Measure Theory, Volume I.* Springer, 2007.

[Bog07b] V. I. Bogachev. *Measure Theory, Volume II.* Springer, 2007.

[Kar75] A. F. Karr. Weak convergence of a sequence of markov chains. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 33:41–48, 1975.

[Lau96] S. L. Lauritzen. *Graphical Models.* Oxford University Press, 1996.

[PB07] R. Pfeifer and J. Bongard. *How the body shapes the way we think.* MIT Press, 2007.

[Pea] J. Pearl. *Causality: Models, Reasoning and Inference.*

[TP] N. Tishby and D. Polani. Information theory of decisions and actions. *In Cutsuridis, V., Hussain, A., and Taylor, J., editors, Perception-Action Cycle: Models, Architecture and Hardware.*

[Uex34] J. Von Uexküll. *A stroll through the worlds on animals and men.* In: Instructive Behavior, ed. K. Lashley. International University Press, 1934.

[ZAD] K. Zahedi, N. Ay, and R. Der. Higher coordination with less control: A result of information maximization in the sensori-motor loop. *Adaptive Behavior.*