The crawling robot uses its HLEC skin to sense its physical state and environment (i.e., proprioception and exteroception). The capacitance of the HLEC changes with pneumatic actuation (Fig. 4B and data S8) and externally applied pressure (Fig. 4, C and D, and data table S9) (*10*). Actuation of the three underlying pneumatic chambers results in capacitance changes (ΔC) of up to 1000% when the chambers are fully inflated. Additionally, each HLEC panel is largely decoupled from the state of the surrounding pneumatic chambers (fig. S4 and data table S11) (*28*). The ability to identify the actuated state of the robot using the capacitive sensor readings enables proprioception. To demonstrate the tactile sensing capabilities of the electronic skin, we pressed each of the HLEC panels on the robot and measured the capacitive response (Fig. 4C). A firm finger press resulted in a ~25% increase in capacitance. The relative capacitance versus applied pressure, ranging from 0.9 to 30.9 kPa, remained nearly constant over a period of 120 hours (Fig. 4D). Arrays of these tactile sensors enable exteroception in soft robotic systems.

An array of three HLEC panels patterned into the three-chambered crawling robot enables eight distinct illuminated states (Fig. 4E). The embedded HLEC remains functional as the robot is actuated through its crawling sequence (Fig. 4F and movie S3). During actuation, the embedded HLEC undergoes stretches of $\lambda_1 = 2.63$ and $\lambda_2 = 2.42$ in the longitudinal (front to rear) and transverse (side to side) directions, respectively, to produce a ~635% increase in the skin's surface area (fig. S5). Similar to the single-panel HLEC (movie S1), the luminescence of the embedded skin increases during actuation as its thickness is decreased.

Integrating these highly stretchable and compliant displays into soft actuators enables two new capabilities in soft electronics: (i) displays that actively change their shape and (ii) robots that actively change their color. Using replica molding, we fabricated a multipixel array of individually addressable HLECs, and we used the same process to monolithically integrate these displays into a soft robot capable of changing posture. The HLEC array imparts both dynamic coloration and the potential for feedback control, which would be useful in epidermal electronics (*31*) and robotics (*32*). Although the luminous efficacy of our HLEC (43.2 mlm W^{-1}) is not as high as that of commercial AC powder electroluminescent devices (~4 lm W^{-1}) (*32*), it can be greatly improved by tuning the materials system and device architecture (such as higher-transmissivity encapsulation layers, reduced thickness, and optimized particle size). For applications requiring higher display resolution, HLECs could be made compatible with photolithography and other microfabrication techniques by using photopolymerizable polymers. These techniques would also allow us to decrease the thickness of the electroluminescent layer, thereby reducing the voltage required to power the HLEC.

REFERENCES AND NOTES

1. A. Barbosa, J. J. Allen, L. Mäthger, R. T. Hanlon, *Proc. R. Soc. London Ser. B* **279**, 84–90 (2012).
2. F. Ilievski, A. D. Mazzeo, R. F. Shepherd, X. Chen, G. M. Whitesides, *Angew. Chem. Int. Ed. Engl.* **50**, 1890–1895 (2011).
3. D. Rus, M. T. Tolley, *Nature* **521**, 467–475 (2015).
4. M. J. Spenko *et al.*, *J. Field Robot.* **25**, 223–242 (2008).
5. E. Kreit *et al.*, *J. R. Soc. Interface* **10**, 20120601 (2012).
6. J. A. Rogers, T. Someya, Y. Huang, *Science* **327**, 1603–1607 (2010).
7. S. A. Morin *et al.*, *Science* **337**, 828–832 (2012).
8. Q. Wang, G. R. Gossweiler, S. L. Craig, X. Zhao, *Nat. Commun.* **5**, 4899 (2014).
9. C. Yu *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 12998–13003 (2014).
10. P. E. Burrows *et al.*, *Displays* **22**, 65–69 (2001).
11. T. H. Han *et al.*, *Nat. Photonics* **6**, 105–110 (2012).
12. T. Sekitani *et al.*, *Nat. Mater.* **8**, 494–499 (2009).
13. M. K. Shin *et al.*, *Adv. Mater.* **22**, 2663–2667 (2010).
14. M. S. White *et al.*, *Nat. Photon.* **7**, 811–816 (2013).
15. L. Hu, H. S. Kim, J. Y. Lee, P. Peumans, Y. Cui, *ACS Nano* **4**, 2955–2963 (2010).
16. J. Liang, L. Li, X. Niu, Z. Yu, Q. Pei, *Nat. Photon.* **7**, 817–824 (2013).
17. T. Someya *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 12321–12325 (2005).
18. K. Takei *et al.*, *Nat. Mater.* **9**, 821–826 (2010).
19. T. Someya *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 9966–9970 (2004).
20. Y. L. Park, B. R. Chen, R. J. Wood, *IEEE Sens. J.* **12**, 2711–2718 (2012).
21. D. P. J. Cotton, I. M. Graz, S. P. Lacour, *IEEE Sens. J.* **9**, 2008–2009 (2009).
22. C. Keplinger *et al.*, *Science* **341**, 984–987 (2013).
23. J. Y. Sun *et al.*, *Nature* **489**, 133–136 (2012).
24. J. Y. Sun, C. Keplinger, G. M. Whitesides, Z. Suo, *Adv. Mater.* **26**, 7608–7614 (2014).
25. A. Kitai, *Luminescent Materials and Applications* (Wiley, West Sussex, UK, 2008), pp. 249–268.
26. R. F. Shepherd, A. A. Stokes, R. M. D. Nunes, G. M. Whitesides, *Adv. Mater.* **25**, 6709–6713 (2013).
27. Y. Bai *et al.*, *Appl. Phys. Lett.* **105**, 151903 (2014).
28. Materials and methods are available as supplementary materials on *Science* Online.
29. G. F. Alfrey, J. B. Taylor, *Proc. Phys. Soc. B* **68**, 775–784 (1955).
30. R. F. Shepherd *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 20400–20403 (2011).
31. D. H. Kim *et al.*, *Science* **333**, 838–843 (2011).
32. J. Kim *et al.*, *Nat. Commun.* **5**, 5747 (2014).

ACKNOWLEDGMENTS

Data reported in the paper are included in the supplementary materials. This work was supported by the Army Research Office (grant no. W911NF-15-1-0464), the Air Force Office of Scientific Research (grant no. FA9550-15-1-0160), the NSF MRSEC program (DMR-1120296), and an NSF Graduate Research Fellowship (grant no. DGE-1144153). The hyperelastic electroluminescent capacitors presented in this work have been filed under a provisional patent application, no. 62/250,172 for Stretchable Electroluminescent Devices. The listed inventors are Chris Larson, Shuo Li, Bryan Peele, Sanlin Robinson, and Robert Shepherd.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/351/6277/1071/suppl/DC1
Materials and Methods
Supplementary Text
Figs. S1 to S7
Table S1
Reference (*33*)
Movies S1 to S3
Data Tables S1 to S11 (single Excel workbook)

5 May 2015; accepted 4 February 2016
10.1126/science.aac5082

HUMAN ALTRUISM

The brain's functional network architecture reveals human motives

Grit Hein,[1*] Yosuke Morishima,[1,2,3] Susanne Leiberg,[1] Sunhae Sul,[4] Ernst Fehr[1]

Goal-directed human behaviors are driven by motives. Motives are, however, purely mental constructs that are not directly observable. Here, we show that the brain's functional network architecture captures information that predicts different motives behind the same altruistic act with high accuracy. In contrast, mere activity in these regions contains no information about motives. Empathy-based altruism is primarily characterized by a positive connectivity from the anterior cingulate cortex (ACC) to the anterior insula (AI), whereas reciprocity-based altruism additionally invokes strong positive connectivity from the AI to the ACC and even stronger positive connectivity from the AI to the ventral striatum. Moreover, predominantly selfish individuals show distinct functional architectures compared to altruists, and they only increase altruistic behavior in response to empathy inductions, but not reciprocity inductions.

The theory of revealed preference (*1*) provides the choice-theoretic foundations for modern economics. In this view, preferences cannot be identified independently of behavior, and motives play no causal role in economists' explanatory toolbox—a view that is in direct contradiction to the neuroeconomic approach (*2*–*4*). In psychology, motives are also considered to be independent drivers of goal-directed human behavior (*5*). Motives are, however, mental constructs that are not directly observable and frequently not even accessible introspectively, meaning that asking people does not provide relevant information about motives (*6*, *7*). Therefore, human motives have been typically inferred from individuals' behavior by assuming that different motives lead to different behaviors.

[1]Laboratory for Social and Neural Systems Research, Department of Economics, University of Zurich, Switzerland. [2]Division of Systems Neuroscience of Psychopathology, Translational Research Center, University Hospital of Psychiatry, University of Bern, Switzerland. [3]Japanese Science and Technology Agency, PRESTO, Japan. [4]Department of Psychology, Pusan National University, Pusan, South Korea.
***Corresponding author. E-mail: grit.hein@econ.uzh.ch (G.H.); ernst.fehr@econ.uzh.ch (E.F.)**

Here we ask whether different motives have a distinct neurophysiological representation that is generalizable across individuals. That is, even if we had no information about individuals' behaviors or if these behaviors would not allow us to make inferences about motives, could we still identify and predict their motives merely on the basis of their functional neural network architecture?

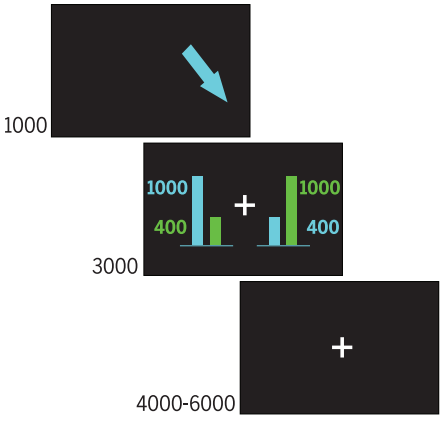

Fig. 1. Example trial in the allocation task in which subjects could allocate money to themselves (indicated by the green bars in the second screen) and one of their partners (indicated by the blue bars in the second screen). Subjects could either maximize their own monetary payoff (selfish decision) or maximize the paired partner's monetary payoff by giving up their own payoff (altruistic decision). They faced several trials during the allocation task that differed in the subject's cost of maximizing the partner's payoff (*21*). In each trial, subjects from the empathy induction group were paired with the empathy partner (empathy condition) or the baseline partner (baseline condition). Likewise, in each trial, subjects from the reciprocity induction group were paired with a reciprocity partner (reciprocity condition) or the baseline partner (baseline condition). The arrow in the first screen indicates the partner with whom the subject was paired in a trial. The choice problems in the empathy, the reciprocity, and the baseline conditions were identical.

We tackled this question in the context of human altruistic decisions (*8*–*14*). Subjects participated in an allocation task in which they could make selfish or altruistic decisions. We studied the role of two key motives for altruistic behaviors—the empathy motive and the reciprocity motive, two important drivers for human altruism (*8*–*14*). We induced these motives experimentally in two different groups of subjects, i.e., subjects were randomly assigned to either the empathy induction group or the reciprocity induction group. After the motive inductions, subjects participated in the allocation task in which they could allocate money to other individuals at a cost to themselves. All subjects faced the same allocation task regardless of the previous motive-induction group. Therefore, their underlying motive cannot be inferred from the mere fact that they behave altruistically. Can we now predict the induced motive solely on the basis of the subject's functional neural architecture?

We used dynamic causal models (DCMs) of functional magnetic resonance imaging data (*15*–*17*) collected during the allocation task and used the estimated DCM parameters to predict subjects' "hidden" motivational state with machine learning—an approach known as generative embedding (*17*). More specifically, the DCM analyses of subjects' brain data during altruistic decision-making gave us information about individuals' network architecture in the different motive conditions. These parameters then provided the "raw" material for our predictions and for the mechanistic insights that follow from our examination (*17*).

Both in the empathy and the reciprocity induction group, subjects were paired with two partners (confederates of the experimenter), who were sitting on either side of the subject. In the empathy induction group, the subject repeatedly observed one of the confederates (the empathy partner) receiving pain shocks in a number of trials, a situation known to elicit an empathic response (*18*, *19*). The reciprocity motive is defined as the desire to reciprocate perceived kindness with a kind behavior (*13*, *14*). Therefore, in the reciprocity induction group, we activated the reciprocity motive by instructing one of the confederates (the reciprocity partner) to give up money in several trials to save the subject from painful shocks (*13*, *14*, *20*). No motives were induced toward the respective second partner (baseline partner), who played the same role in both the empathy- and the reciprocity induction group [for details of motive induction, see supplementary materials (*21*)]. It is important to stress that the subjects received painful shocks not only in the reciprocity induction group but also in the empathy induction group, and that the number of painful shocks was identical across conditions (*21*). This feature has two advantages. First, by equalizing the shock frequency across conditions, we can be sure that the two motive inductions contain the same number of aversive events. Second, the application of painful shocks to the subject in the empathy condition is likely to enhance the ability to empathize with the empathy partner because subjects know how the shock feels. Finally, to assess the success of the motive inductions, the subjects also completed emotion ratings where they indicated in each trial how they felt.

In the allocation task, subjects were in the scanner. In each trial, they allocated money between themselves and one of the partners (Fig. 1). They could choose between maximizing the other person's monetary payoff by reducing their own monetary payoff (altruistic behavior), or maximizing their own payoff (selfish behavior) at a cost to the partner. Depending on the type of partner subjects faced in the allocation task, there were three conditions—the empathy condition, the reciprocity condition, and the baseline condition. Because neither reciprocity nor empathy motives were induced in the baseline condition

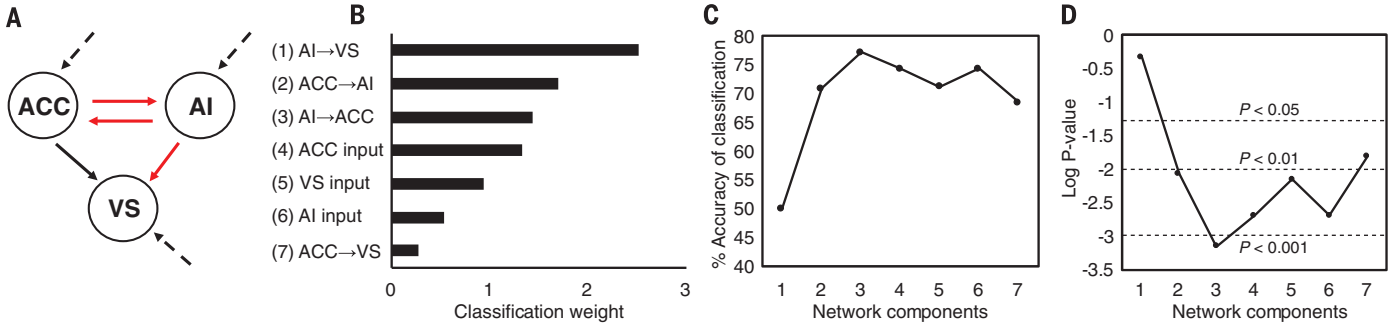

Fig. 2. DCM-based classification of different motives that resulted in the same altruistic decision. (**A**) Dynamic causal model (DCM) that was used to investigate the impact of the empathy and the reciprocity motive on brain connectivity during altruistic decisions. Straight arrows indicate interregional connections (four components), and dashed arrows represent stimulus inputs from outside of the network (three components). (**B**) Discriminative components. The figure visualizes the weights with which each of the seven network components contributed to the classification of the empathy and reciprocity motive. (**C**) Accuracy for each level of the stepwise classification, based on the order of classification weights in (B). (**D**) Log *P*-value of stepwise classification, based on the order of classification weights in (B). AI, anterior insula; ACC, anterior cingulate cortex; VS, ventral striatum. Network components: 1 = AI→ VS; 2 = plus ACC→AI; 3 = plus AI→ACC; 4 = plus ACC input; 5 = plus VS input; 6 = plus AI input; 7 = plus ACC→VS. The highest accuracy is reached when only the first three of the seven connectivity and input parameters are used to classify the empathy and the reciprocity motive [(A), highlighted in red].

(i.e., vis-à-vis the baseline partner), the behavior in this condition measures subjects' "raw" or "homegrown" unconditional altruism. This altruism is unconditional in the sense that the baseline partner did not do anything during the motive induction on which the subjects could condition their behavior during the allocation task. Homegrown unconditional altruism can play a role in all three conditions, but the induced empathy motive plays an additional role in the empathy condition and the reciprocity motive plays an additional role in the reciprocity condition.

Empathy means that if an individual observes someone else in pain, the individual also feels or "shares" that pain (*18*). Stronger empathizing during the empathy induction should therefore induce worse feelings (due to "shared" pain), which should then lead to a stronger empathy motive for altruistic behavior. We indeed observed that the worse a subject felt when seeing the empathy partner in pain, the more frequent the altruistic decision toward this person in the subsequent allocation task [$r(18) = -0.51$, $P = 0.03$]. Because the reciprocity partner is willing to incur cost to remove painful shocks, the partner is likely to be perceived as kind, which should result in positive feelings and higher frequencies of altruistic decisions toward the reciprocity partner. We found indeed that the better a subject felt after the reciprocity partner's decision during motive induction, the more altruistic the decisions toward this person in the subsequent allocation task [$r(16) = 0.57$, $P = 0.021$]. Reciprocity induction also induced significantly higher likability ratings—which were collected after scanning ("How much do you like the other person"? 9 = very much to 1 = not at all)—for the reciprocity partner as compared to the baseline partner [$t(15) = 3.24$, $P = 0.005$]. The overall frequency of altruistic decisions toward the empathy partner and the reciprocity partner was significantly higher than toward the baseline partner [motive induction versus baseline, $F(1,32) = 12.5$, $P = 0.001$; empathy versus baseline, $t(17) = 2.5$, $P = 0.022$; reciprocity versus baseline, $t(15) = 2.7$, $P = 0.017$]. There was, however, no significant difference in the increase of altruistic decisions between the two motive inductions [empathy versus baseline compared to reciprocity versus baseline, $F(1,32) = 0.2$, $P = 0.64$]. In addition, the results of a Bayesian analysis show that the null hypothesis that the two motive inductions cause the same increase in altruism relative to baseline is more than five times more likely (84% versus 16%) than the hypothesis of a differential increase (*21*).

Our imaging analyses focused on altruistic decisions during the allocation task. First we used a conventional general linear model (GLM) analysis for a whole-brain search of regions with significantly different activity during altruistic decisions driven by empathy versus altruistic decisions driven by reciprocity. There were no significant differences even at the very liberal threshold of $P_{uncorrected} < 0.05$. This suggests that empathy-driven and reciprocity-driven altruism activate similar brain regions. We therefore compared brain activations under altruistic decisions in both motive-induction conditions with the activations involved in altruistic decisions in the baseline condition. The results revealed a network consisting of left anterior insula (AI), left ventral striatum (VS), and anterior cingulate cortex (ACC) [$P < 0.05$; family-wise error (FWE) corrected; see table S1 for details]—regions that were reported by previous studies on the reciprocity (*22*–*24*) and the empathy motive (*18*, *19*, *25*).

In a second step, we determined the pattern of neural connectivity within this network for each subject, using DCM. We extracted the time series of activations during altruistic decisions toward the empathy partner, the reciprocity partner, and the respective baseline partner from individual regions of interest (ROIs) in left AI, left VS, and ACC (table S2).

The DCM analyses were based on the anatomical model shown in Fig. 2A, which is characterized by seven components: three inputs (dashed arrows) and four interregional connectivities (solid arrows). Because most VS projections target the cortex via the thalamus (*26*, *27*), ascending projections from VS were not included in the anatomical model. We used Bayesian model averaging (*28*) to determine the DCM parameters for the seven components for each subject. For a subject of the empathy group, for example, we calculated (i) how activation in ACC, AI, and VS during altruistic decisions is changed as a result of impulses from outside the network under the empathy and under the baseline conditions (dashed arrows in Fig. 2A), and (ii) how activation in every (target) region is changed by the level of activation in the other regions during altruistic decisions under the empathy and the baseline conditions (effective connectivity, solid arrows in Fig. 2A). The same computations were also done for subjects of the reciprocity group. We thus obtained 14 DCM parameters per subject, seven for altruistic decisions under the respective motive-induction condition (empathy or reciprocity), and seven for altruistic decisions under the baseline condition. We then subtracted the individual DCM parameters of the baseline condition from the DCM parameters of the motive-induction condition. The resulting seven Δ-DCM parameters reflect the individual pattern of neural connectivity specific for altruistic decisions driven by empathy or reciprocity.

In a third step, we submitted these individual Δ-DCM parameters to a classification algorithm (support vector machine, SVM) to test if the individual patterns of brain connectivity can be used to detect the specific motive for altruism that the experimenter induced (*17*). This approach

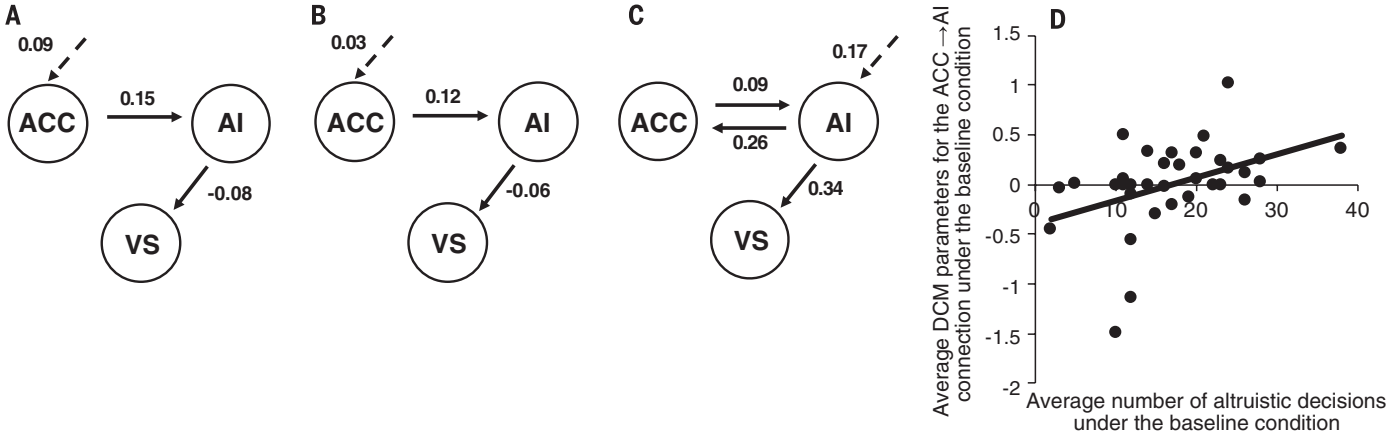

Fig. 3. Average best models and mean DCM parameters. Best models and mean DCM parameters for (**A**) Altruistic decisions driven by the empathy motive, (**B**) altruistic decisions driven by "homegrown" unconditional altruism in the baseline condition, and (**C**) altruistic decisions driven by the reciprocity motive. For visualization purposes, the baseline data for subjects from the empathy and the reciprocity group were pooled in (B), because there were no significant differences in the corresponding best-model parameters [ACC→AI, $t(1,32) = -0.57$, $P = 0.58$; AI→VS, $t(1,32) = -1.03$, $P = 0.31$; ACC input, $t(1,32) = 0.55$, $P = 0.59$]. In (A) to (C), dashed arrows indicate direct inputs to an area; solid arrows indicate directed interregional connectivities. AI, anterior insula; ACC, anterior cingulate cortex; VS, ventral striatum. (**D**) Correlation between the individual parameters of the ACC→AI connection under baseline conditions and the individual average number of altruistic decisions toward the baseline partner.

significantly predicted the subjects' induced altruistic motives, i.e., whether subjects went through the empathy or the reciprocity induction condition (classification accuracy = 68.4%, $P = 0.016$). The computation of classification accuracy with the associated P-value is based on the posterior probability of balanced accuracy as in (*17*). We also tested if the two motives can be classified on the basis of conventional functional activations. We extracted the corresponding beta values from the traditional GLM analysis (empathy versus baseline condition; reciprocity versus baseline condition) from the same ROIs that were used for the DCM analyses and submitted them to the classification algorithm. This analysis did not yield classification accuracy significantly above chance (classification accuracy = 55.2%, $P = 0.3$). It was also not possible to classify the motives based on the increases in the frequency of altruistic decisions in the two motive-induction conditions relative to the baseline condition [classification accuracy = 41%, $P = 0.93$; see also (*29*)].

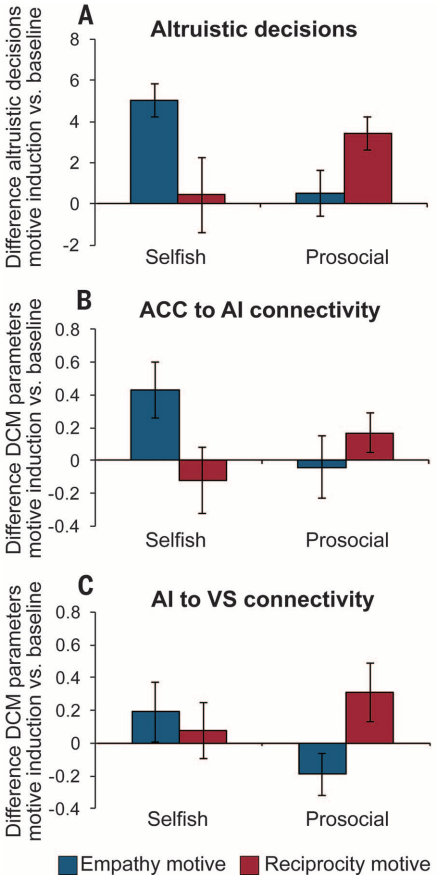

Fig. 4. Differential impact of the empathy and the reciprocity motive on individuals with predominantly selfish and prosocial preferences. (**A**) Change in the frequency of altruistic decisions after the empathy and the reciprocity induction, compared to baseline, in selfish and prosocial individuals. The empathy induction only increases the frequency of altruistic decisions in more selfish subjects, whereas the reciprocity induction only increases the frequency of altruistic decisions in subjects with more prosocial preferences. (**B**) Change in brain connectivity from the ACC to the AI after the empathy and reciprocity induction, compared to baseline, in selfish and prosocial individuals. (**C**) Change in brain connectivity from the AI to the VS after the empathy and reciprocity induction, compared to baseline, in selfish and prosocial individuals. Error bars indicate SEM.

To better understand which DCM parameters jointly enabled the distinction between the empathy and the reciprocity motive, we calculated the average classification weights with which each of the seven different network components contributed to the classification (Fig. 2B), and the stepwise classification performance of the sequential combination of network components (Fig. 2, C and D). The strongest classification weight was attached to the connection from AI to VS, followed by the directional connection from the cingulate cortex to the insular cortex and vice versa (see red arrows in Fig. 2A). The combination of these three network components was sufficient for classifying the induced motives behind subjects' altruistic decisions with an accuracy of 77% ($P = 0.0007$; Fig. 2, C and D).

To identify the motive-specific functional network architectures, we selected the best average network for altruistic decisions under the empathy, reciprocity, and baseline conditions. The selection procedure was based on 28 different models (fig. S3). To define this model space, we used the criteria of anatomical plausibility, network coherence, and functional plausibility [see methods section in supplementary materials (*21*) for details]. We used random-effect Bayesian model selection (*30*) to select the models whose structure and patterns of inputs and effective connectivity fit best with the neural processes evoked by altruistic decisions under the different conditions (*21*).

Figure 3 shows the best models of neural connectivity during empathy-driven, reciprocity-driven, and "homegrown" altruism in the baseline condition (see figs. S3 and S4 for details). A comparison between Fig. 3, A and B, shows a marked similarity in the functional network architecture of empathy-driven and homegrown altruism. Both models show a positive connectivity between ACC and AI, and a slightly negative connectivity between AI and VS, with no significant differences in the respective DCM parameters [ACC to AI, $t(1,17) = -0.24$, $P = 0.8$; AI to VS, $t(1,17) = -0.66$, $P = 0.52$; ACC input, $t(1,17) = 0.89$, $P = 0.34$]. This contrasts sharply with the best model for reciprocity-driven altruism (Fig. 3C). In this model there is, first, a strong bidirectional projection between AI and ACC. Second, there is a strong positive connectivity between AI and VS, which significantly differs from the negative AI-to-VS connectivity under the baseline [$t(1,15) = 2.8$, $P = 0.015$] and empathy conditions [$t(1,32) = 2.91$, $P = 0.006$; $\alpha_{(\text{Bonferroni-corrected})} < 0.025$]. There were no significant differences with regard to the connectivity from ACC to AI [reciprocity versus baseline, $t(1,15) = 0.03$, $P = 0.9$; reciprocity versus empathy, $t(1,32) = -0.28$, $P = 0.78$].

These findings also raise the question whether the differential network components present in the reciprocity condition are related to particular psychological features. Inspired by previous evidence that has linked the social evaluation of other individuals to activation in ventral striatal regions (*31*, *32*), we correlated subjects' likability ratings of the reciprocity partner with the individual strengths of the AI→VS connectivity. There was a significant positive correlation [$r(16) = 0.59$, $P = 0.016$], whereas no such correlation was present in the empathy and the baseline conditions [empathy condition: $r(18) = -0.29$, $P = 0.25$; baseline condition: $r(16) = 0.001$, $P = 0.99$].

There is, however, also a common network feature—the ACC→AI connectivity—that is present in all three conditions. We thus hypothesized that this component of the network might reflect basic prosocial motivation. To test this conjecture, we correlated individuals' frequency of altruistic decisions toward their baseline partner with their DCM parameters of the ACC→AI connection. The results show a significant positive correlation [$r(34) = 0.4$, $P = 0.017$] (Fig. 3D). The stronger a person's connectivity from ACC to AI, the higher the baseline level of altruism. Indeed, if we divide our sample of subjects in prosocial and selfish individuals on the basis of a median split in the frequency of altruistic decisions in the baseline condition, we find that prosocial individuals display a positive connectivity from ACC to AI of 0.16, whereas selfish individuals show a negative connectivity of −0.17—a difference that is highly significant [$F(1,32) = 9.49$, $P = 0.004$].

If selfish and prosocial subjects display different network architectures, they may also respond differently to the empathy and the reciprocity induction. We sorted the selfish and prosocial individuals by the respective motive-induction condition; this resulted in nine prosocial subjects in both the empathy and the reciprocity condition, nine selfish subjects in the empathy condition, and seven selfish subjects in the reciprocity condition. Figure 4A shows that the induction of the empathy motive significantly increased altruistic decisions in selfish individuals, whereas there was no such effect after the induction of reciprocity. In contrast, the induction of reciprocity resulted in a further enhancement of altruistic behavior in prosocial individuals, whereas the empathy induction had no effect on these individuals. Thus, the two types of subjects respond very differently to the two motive inductions [motive induction × individual type, $F(1,30) = 8.8$, $P = 0.006$].

We next compared the differences in brain connectivity between the motive-induction and the baseline conditions. Empathy induction increased effective connectivity from ACC to AI in selfish subjects, while there was no such effect after the selfish subjects received a reciprocity induction and in prosocial individuals [motive induction × individual type, $F(1,30) = 4.1$, $P = 0.05$] (Fig. 4B). In contrast, the reciprocity induction led to an enhancement of neural connectivity from AI to VS in prosocial individuals, which is not the case

after the empathy induction and in selfish types [motive induction × individual type, $F(1,30) = 4.9$, $P = 0.034$] (Fig. 4C).

Motives are purely mental constructs that are not directly observable. Here we show, however, that distinct motives have a distinct neurophysiological representation in the brain. Although the empathy and the reciprocity motive increase the frequency of altruistic acts by the same amount relative to the baseline condition, they are associated with different patterns of brain connectivity that enabled us to predict the different motives with relatively high accuracy. We predicted each subject's induced motive with a classifier whose parameters were not influenced by that subject's brain data (nor by that subject's behavioral data). Instead, the parameters of the classifier were solely informed by other subjects' brain data. This means that the motive-specific brain connectivity patterns are generalizable across subjects. The distinct and across-subject–generalizable neural representation of the different motives thus provides evidence for a distinct neurophysiological existence of motives.

The findings also provide mechanistic insights into the neural underpinnings of important altruistic motives and how motive inductions change the underlying neural network. In particular, predominantly selfish individuals were characterized by a low or even negative connectivity from ACC→AI in the baseline condition, whereas predominantly prosocial individuals displayed a positive connectivity between these regions. However, when we induce the empathy motive, the selfish, but not the prosocial, types become more altruistic and show a substantial increase in ACC→AI connectivity. Thus, after the empathy induction, selfish individuals resemble "homegrown" unconditional altruists in terms of both brain connectivity and altruistic behavior. This contrasts with the effect of inducing the reciprocity motive, which renders the prosocial, but not the selfish, types more altruistic and increases only the prosocial types' AI→VS connectivity.

We obtain these mechanistic insights because the inputs into the support vector machine are not merely brain activations but small brain models of how relevant brain regions interact with each other (i.e., functional neural architectures). Thus, by correctly predicting the induced motives, we simultaneously determine those mechanistic models of brain interaction that best predict the motives. And it is these models that deliver the mechanistic insights into brain function and how changes in brain function relate to behavioral changes due to motive inductions. Our study, therefore, also demonstrates how "mere prediction" and "insights into the mechanisms" that underlie psychological concepts (such as motives) can be simultaneously achieved if functional neural architectures are the inputs for the prediction.

REFERENCES AND NOTES

1. P. A. Samuelson, *Economica* **15**, 243–253 (1948).
2. P. W. Glimcher, A. Rustichini, *Science* **306**, 447–452 (2004).
3. A. Rangel, C. Camerer, P. R. Montague, *Natl. Rev.* **9**, 545–556 (2008).
4. E. Fehr, A. Rangel, *J. Econ. Perspect.* **25**, 3–30 (2011).
5. K. Lewin, *Field Theory in Social Science* (Harper, New York, 1951).
6. R. B. Zajonc, *Am. Psychol.* **35**, 151–175 (1980).
7. R. E. Nisbett, T. D. Wilson, *Psychol. Rev.* **84**, 231–259 (1977).
8. C. D. Batson, *The Altruism Question: Toward a Social-Psychological Answer* (Erlbaum, Hillsdale, NJ, 1991).
9. E. Fehr, U. Fischbacher, *Nature* **425**, 785–791 (2003).
10. J. Moll *et al*., *Proc. Natl. Acad. Sci. U.S.A.* **103**, 15623–15628 (2006).
11. W. T. Harbaugh, U. Mayr, D. R. Burghart, *Science* **316**, 1622–1625 (2007).
12. Y. Morishima, D. Schunk, A. Bruhin, C. C. Ruff, E. Fehr, *Neuron* **75**, 73–79 (2012).
13. M. Rabin, *Am. Econ. Rev.* **83**, 1281–1302 (1993).
14. A. Falk, U. Fischbacher, *Games Econ. Behav.* **54**, 293–315 (2006).
15. K. J. Friston, L. Harrison, W. Penny, *Neuroimage* **19**, 1273–1302 (2003).
16. K. E. Stephan *et al*., *Neuroimage* **42**, 649–662 (2008).
17. K. H. Brodersen *et al*., *PLOS Comput. Biol.* **7**, e1002079 (2011).
18. T. Singer *et al*., *Science* **303**, 1157–1162 (2004).
19. G. Hein, G. Silani, K. Preuschoff, C. D. Batson, T. Singer, *Neuron* **68**, 149–160 (2010).
20. R. A. Emmons, M. E. McCullough, *The Psychology of Gratitude* (Oxford Univ. Press, 2004).
21. Methods and materials, supplementary analyses, supplementary figures, and supplementary tables are available as supporting material on *Science* Online.
22. D. Tomlin *et al*., *Science* **312**, 1047–1050 (2006).
23. J. K. Rilling *et al*., *Neuropsychologia* **46**, 1256–1266 (2008).
24. K. L. Phan, C. S. Sripada, M. Angstadt, K. McCabe, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 13099–13104 (2010).
25. C. Lamm, J. Decety, T. Singer, *Neuroimage* **54**, 2492–2502 (2011).
26. S. N. Haber, B. Knutson, *Neuropsychopharmacol* **35**, 4–26 (2010).
27. A. Parent, L. N. Hazrati, *Brain Res. Brain Res. Rev.* **20**, 91–127 (1995).
28. J. A. Hoeting, D. Madigan, A. E. Raftery, C. T. Volinsky, *Stat. Sci.* **14**, 382 (1999).
29. If we use both the number of altruistic decisions in the baseline condition and the increase in the frequency of altruistic decisions in the motive-induction conditions, the behavioral classification becomes marginally significant (classification accuracy of 64.2%, $P = 0.051$). However, if we perform the same classification analysis with connectivity data—i.e., in addition to the Δ-DCM parameters we also use the level of the DCM parameters in the baseline condition for classification purposes—the classification accuracy increases even to 83%, $P = 0.00004$. Thus, the classification based on brain connectivity data clearly outperforms the behavior-based classification (see also supplementary materials).
30. K. E. Stephan, W. D. Penny, J. Daunizeau, R. J. Moran, K. J. Friston, *Neuroimage* **46**, 1004–1017 (2009).
31. M. R. Delgado, R. H. Frank, E. A. Phelps, *Nat. Neurosci.* **8**, 1611–1618 (2005).
32. A. Tusche, T. Kahnt, D. Wisniewski, J. D. Haynes, *Neuroimage* **72**, 174–182 (2013).

ACKNOWLEDGMENTS

We thank K. E. Stephan for useful comments on early versions of the manuscript and K. Treiber and S. Klein for assistance with data collection. We also thank two anonymous referees for their helpful comments. The paper is part of the advanced European Research Council grant on the "Foundations of Economic Preferences and the Synergia grant of the Swiss National Science Foundation on the neuroeconomics of value-based decision making (CRSII3_141965)." The data are stored on the server of the Laboratory for Social and Neural Systems Research at the University of Zurich.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/351/6277/1074/suppl/DC1
Materials and Methods
Figs. S1 to S4
Tables S1 and S2
References (*33*–*37*)

15 June 2015; accepted 22 January 2016
10.1126/science.aac7992

EBOLA VIRUS

Isolation of potent neutralizing antibodies from a survivor of the 2014 Ebola virus outbreak

Zachary A. Bornholdt,[1] Hannah L. Turner,[2] Charles D. Murin,[1,2] Wen Li,[3] Devin Sok,[1] Colby A. Souders,[4] Ashley E. Piper,[5] Arthur Goff,[5] Joshua D. Shamblin,[5] Suzanne E. Wollen,[5] Thomas R. Sprague,[5] Marnie L. Fusco,[1] Kathleen B. J. Pommert,[1] Lisa A. Cavacini,[4] Heidi L. Smith,[4] Mark Klempner,[4] Keith A. Reimann,[4] Eric Krauland,[3] Tillman U. Gerngross,[3] Karl D. Wittrup,[3] Erica Ollmann Saphire,[1] Dennis R. Burton,[1,6] Pamela J. Glass,[5] Andrew B. Ward,[2] Laura M. Walker[3]*

Antibodies targeting the Ebola virus surface glycoprotein (EBOV GP) are implicated in protection against lethal disease, but the characteristics of the human antibody response to EBOV GP remain poorly understood. We isolated and characterized 349 GP-specific monoclonal antibodies (mAbs) from the peripheral B cells of a convalescent donor who survived the 2014 EBOV Zaire outbreak. Remarkably, 77% of the mAbs neutralize live EBOV, and several mAbs exhibit unprecedented potency. Structures of selected mAbs in complex with GP reveal a site of vulnerability located in the GP stalk region proximal to the viral membrane. Neutralizing antibodies targeting this site show potent therapeutic efficacy against lethal EBOV challenge in mice. The results provide a framework for the design of new EBOV vaccine candidates and immunotherapies.

In recent years, Ebola virus (EBOV) outbreaks have increased in frequency, duration, and geographical spread, underscoring the need for pre- and post-exposure treatments (*1*). The membrane-anchored EBOV glycoprotein (GP) trimer is the sole known target for protective antibodies and is currently the primary target for antiviral vaccines and therapies (*2*, *3*). A small number of protective monoclonal antibodies (mAbs) to GP have been isolated from immunized mice,

The brain's functional network architecture reveals human motives

Grit Hein *et al.*

Science **351**, 1074 (2016);
DOI: 10.1126/science.aac7992