

Q3 Stroke Scenario There are more than 100,000 strokes in the UK each year – that is around one stroke every five minutes. About 11% of patients die immediately or within a few weeks as a result of the stroke, making stroke the fourth biggest killer in the UK. Almost two thirds of stroke survivors leave hospital with a disability. Rapid and accurate diagnosis of stroke greatly increases chances of survival and recovery of the patient. This is highly specialised work which ideally should be done by neuroradiologists with many years of training and experience. However, these experts are not available in each hospital, 24 hours a day, 7 days a week, and in practice diagnosis is often done by non-specialist emergency medicine doctors. As diagnostic data are accumulated from previous stroke patients, automated decision systems could provide stroke diagnosis that is fast, and always available in each hospital. There are three automated decision systems for the NHS to choose from – system A, system B, and system C. Each system uses information about a patient's acute symptoms (for instance paralysis and loss of speech), their medical history, and neuroradiological images (such as CT-scans of the brain) to identify patterns that indicate whether he or she has had a stroke; the type of stroke; its location; and its severity.

System A – Expert System This system uses an algorithm that was developed with help from experienced neurologists and neuroradiologists, and aims to follow the same reasoning as they would do. In practice it does not reach the same level of accuracy as they would, but the algorithm is completely transparent in the way it reaches its conclusions: for each individual case it can provide specific rules that were applied to reach a conclusion. It has an overall accuracy rate of 75%, which is comparable to what most emergency medicine doctors would achieve. This means that in 25% of cases, someone might be classified as having a stroke while they were not or vice versa, or the type, location, and severity of the stroke might be misjudged.

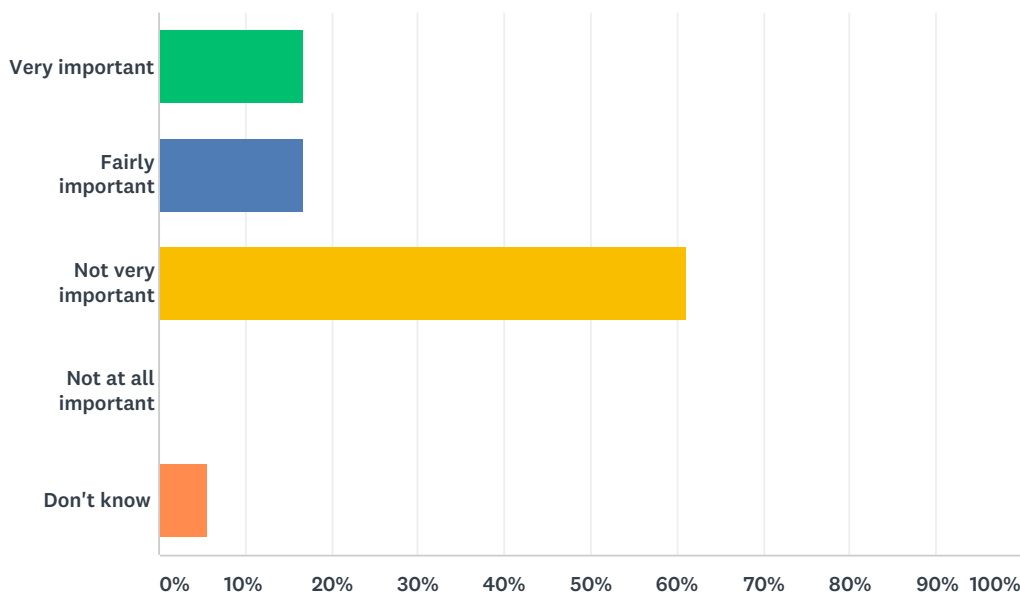
System B – Conventional machine learning This system uses an algorithm that was established through machine learning from a large set of patient data, collected at English hospitals. This algorithm reaches (human) expert level performance, but it is not very transparent in the way it reaches its conclusions: it can only tell us which features, in general, are important and which are not. It has an overall accuracy rate of 85%. This means that in 15% of cases, someone might be classified as having a stroke while they were not or vice versa, or the type, location, and severity of the stroke might

be misjudged. System C – Deep Learning This system uses advanced AI derived from the same set of patient data as System B. However it has “taught itself” from the data which features were best able to distinguish strokes from non-strokes, and best able to distinguish different types of stroke, their location, and their severity. This algorithm is not transparent in the way it reaches conclusions: it is unable to provide any explanation that is understandable by humans. However it has an overall accuracy rate of 95%, which is better than human experts perform. This means that in 5% of cases, someone might be classified as having a stroke while they were not or vice versa, or the type, location, and severity of the stroke might be misjudged.

System A: 75% (A&E doctor’s level) accuracy, full explanation
 System B: 85% (human expert level) accuracy, partial explanation
 System C: 95% (beyond human level) accuracy, no explanation

Q1. How important is it for a patient to receive an explanation of an automated decision about stroke diagnosis?

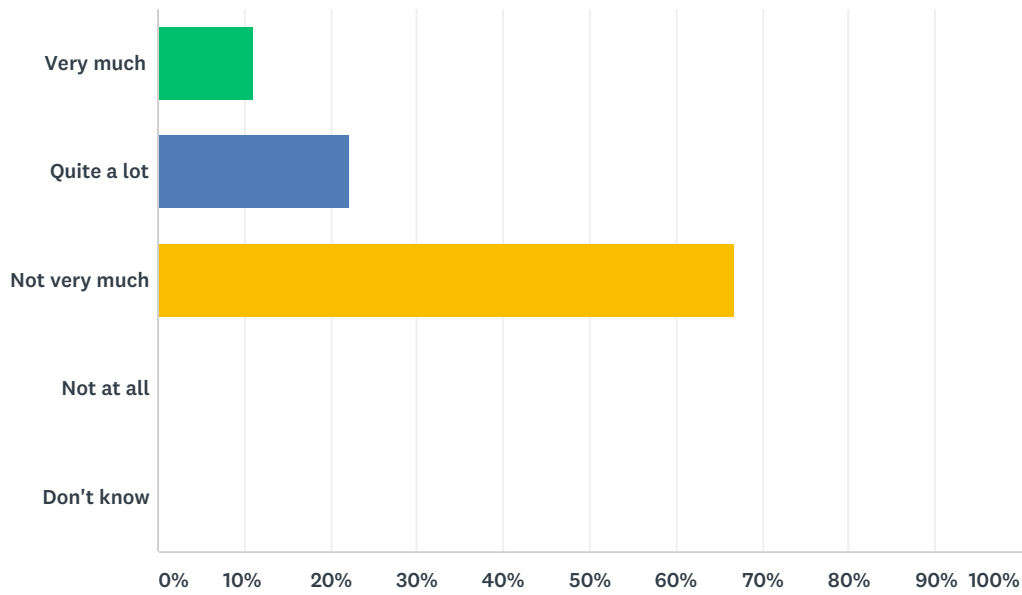
Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------------|-----------|-----------|
| Very important | 16.67% | 3 |
| Fairly important | 16.67% | 3 |
| Not very important | 61.11% | 11 |
| Not at all important | 0.00% | 0 |
| Don't know | 5.56% | 1 |
| TOTAL | | 18 |

Q4 Q2 If system C was chosen by the NHS, almost no explanation would be provided. How much does this matter?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|-----------|
| Very much | 11.11% | 2 |
| Quite a lot | 22.22% | 4 |
| Not very much | 66.67% | 12 |
| Not at all | 0.00% | 0 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

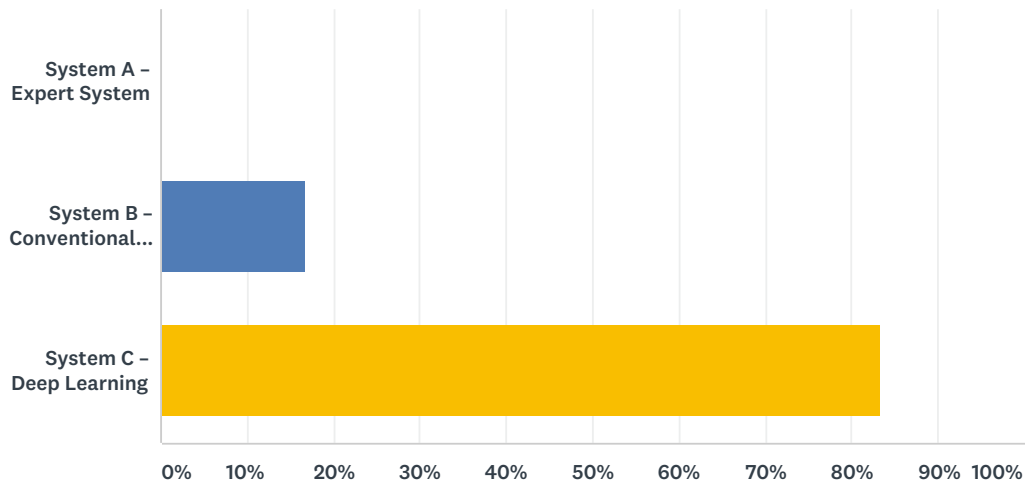
Q5 Why (up to three reasons)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Reason 1 | 100.00% | 18 |
| Reason 2 | 83.33% | 15 |
| Reason 3 | 66.67% | 12 |

Q6 Q3 Which automated decision system do you think the NHS should choose?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES |
|--|-----------|
| System A – Expert System | 0.00% 0 |
| System B – Conventional machine learning | 16.67% 3 |
| System C – Deep Learning | 83.33% 15 |
| TOTAL | 18 |

Q7 Explain the factors affecting your choice (up to 3)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Factor 1 | 100.00% | 18 |
| Factor 2 | 77.78% | 14 |
| Factor 3 | 55.56% | 10 |

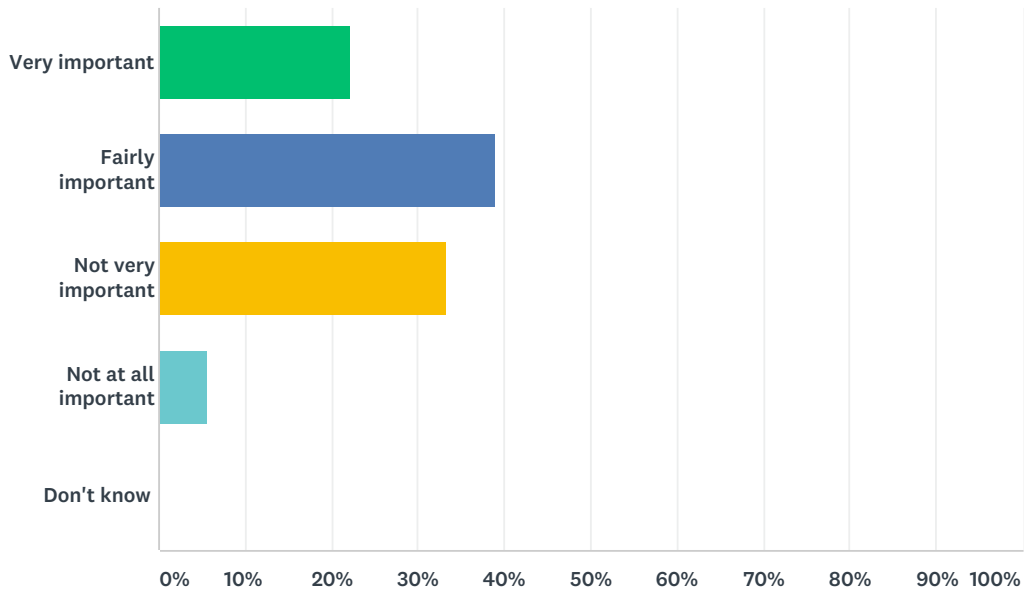
Q3 There are 3 automated decision systems for the organisation to choose from:

- System A – Expert System This system uses an algorithm that was developed with help from experienced recruitment officers, and aims to follow the same reasoning as they would do. In practice it does not reach the same level of accuracy as they would, but the algorithm is completely transparent in the way it reaches its conclusions: for each individual case it can provide specific rules that were applied to reach a conclusion. When tested on existing data about recruitment, this system was shown to have an overall accuracy rate of 75%. This means that 25% of the time its predictions were incorrect (e.g. predicting that an applicant would be unlikely to become a high-performing employee when in reality they did, or vice versa). The accuracy of this system is comparable to that of a typical recruitment officer.
- System B – Conventional machine learning This system uses an algorithm that was established through machine learning from a large set of recruitment data, collected by the organisation. This algorithm achieves (human) expert level performance, but it is not very transparent in the way it reaches its conclusions: it can only tell us which features, in general, are important and which are not. When tested on existing data about recruitment, this system was shown to have an overall accuracy rate of 85%. This means that 15% of the time its predictions were incorrect (e.g. predicting that an applicant would be unlikely to become a high-performing employee when in reality they did, or vice versa). The accuracy of this system is comparable to that of a very experienced recruitment officer.
- System C – Deep Learning This system uses advanced AI, derived from the same set of data as System B. However it has “taught itself” from the data. This algorithm is not transparent in the way it reaches conclusions: it is unable to provide any explanation that is understandable by humans. However it has an overall accuracy rate of 95%, which is better than human experts perform. This means that 5% of the time its predictions were incorrect (e.g. predicting that an applicant would be unlikely to become a high performing employee when in reality they did, or vice versa).

System A: 75% (recruitment officer level) accuracy, full explanation
System B: 85% (human expert level) accuracy, partial explanation
System C: 95% (beyond human level) accuracy, no explanation

Q1. How important is it for an applicant to receive an explanation of an automated decision about accepting / rejecting a job application?

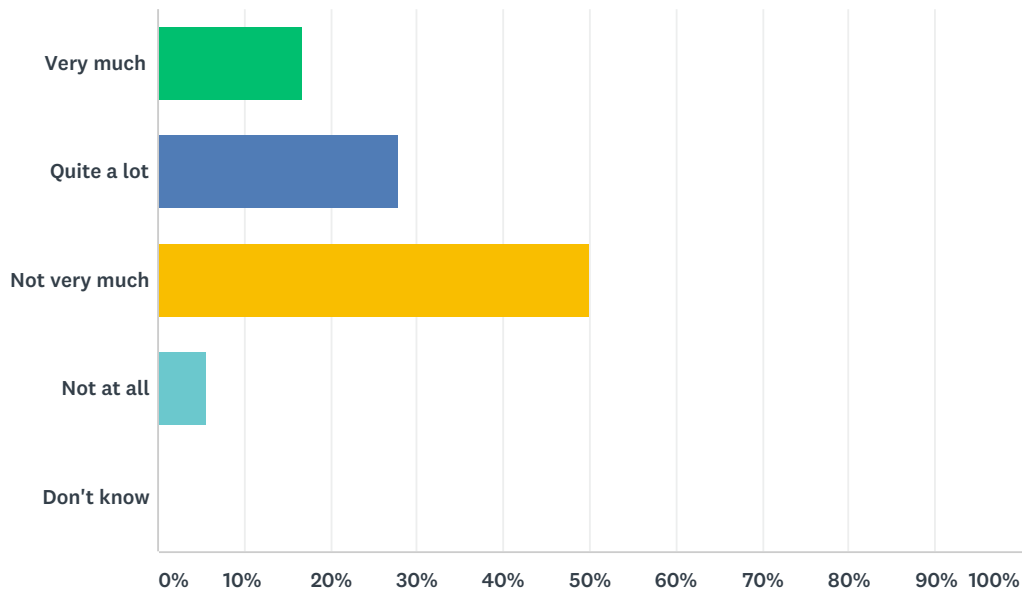
Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------------|-----------|-----------|
| Very important | 22.22% | 4 |
| Fairly important | 38.89% | 7 |
| Not very important | 33.33% | 6 |
| Not at all important | 5.56% | 1 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

Q4 Q2 If system C was chosen by the organisation, almost no explanation would be provided. How much does this matter?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|-----------|
| Very much | 16.67% | 3 |
| Quite a lot | 27.78% | 5 |
| Not very much | 50.00% | 9 |
| Not at all | 5.56% | 1 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

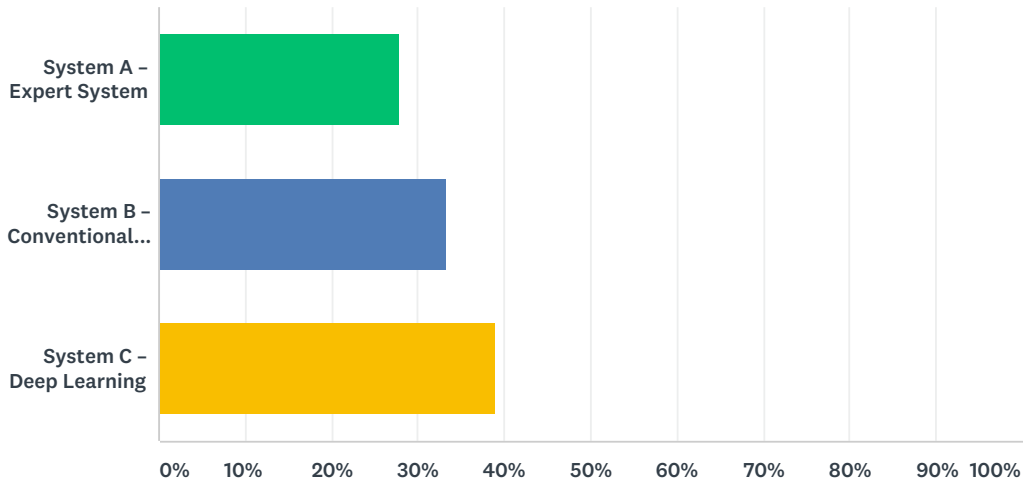
Q5 Why (up to three reasons)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Reason 1 | 100.00% | 18 |
| Reason 2 | 61.11% | 11 |
| Reason 3 | 38.89% | 7 |

Q6 Q3 Which automated decision system do you think the company should choose?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|--|-----------|----|
| System A – Expert System | 27.78% | 5 |
| System B – Conventional machine learning | 33.33% | 6 |
| System C – Deep Learning | 38.89% | 7 |
| TOTAL | | 18 |

Q7 Explain the factors affecting your choice (up to 3)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Factor 1 | 100.00% | 18 |
| Factor 2 | 66.67% | 12 |
| Factor 3 | 33.33% | 6 |

Q3 It is hoped that with this system, a larger number of transplanted kidneys will survive longer. There are three automated decision systems to choose from – system A, system B, and system C.

System A – Expert System This system uses an algorithm that was developed with help from experienced kidney doctors, and aims to follow the same reasoning as they would do. In practice it does not reach the same level of accuracy as they would, but the algorithm is completely transparent in the way it reaches its conclusions: for each individual case it can provide specific rules that were applied to reach a conclusion. It has an overall accuracy rate of 75%, which is a little lower than what is currently achieved in practice across the NHS (and lower than that achieved by the top specialists). This means that 25% of the time its predictions were incorrect (e.g. predicting that the kidney would last at least 5 years for the selected patient when in reality it didn't).

System B – Conventional machine learning This system uses an algorithm that was established through machine learning from a large set of patient data, collected at English hospitals.. This algorithm achieves (human) expert level performance, but it is not very transparent in the way it reaches its conclusions: it can only tell us which features, in general, are important and which are not. It has an overall accuracy rate of 85%. This means that 15% of the time its predictions were incorrect (e.g. predicting that the kidney would last at least 5 years for the selected patient when in reality it didn't).

System C – Deep Learning This system uses advanced AI, derived from the same set of patient data as System B. However it has “taught itself” from the data which features were best able to distinguish successful matches from non-successful matches. This algorithm is not transparent in the way it reaches conclusions: it is unable to provide any explanation that is understandable by humans. However it has an overall accuracy rate of 95%, which is better than human experts perform. This means that 5% of the time its predictions were incorrect (e.g. predicting that the kidney would last at least 5 years for the selected patient when in reality it didn't).

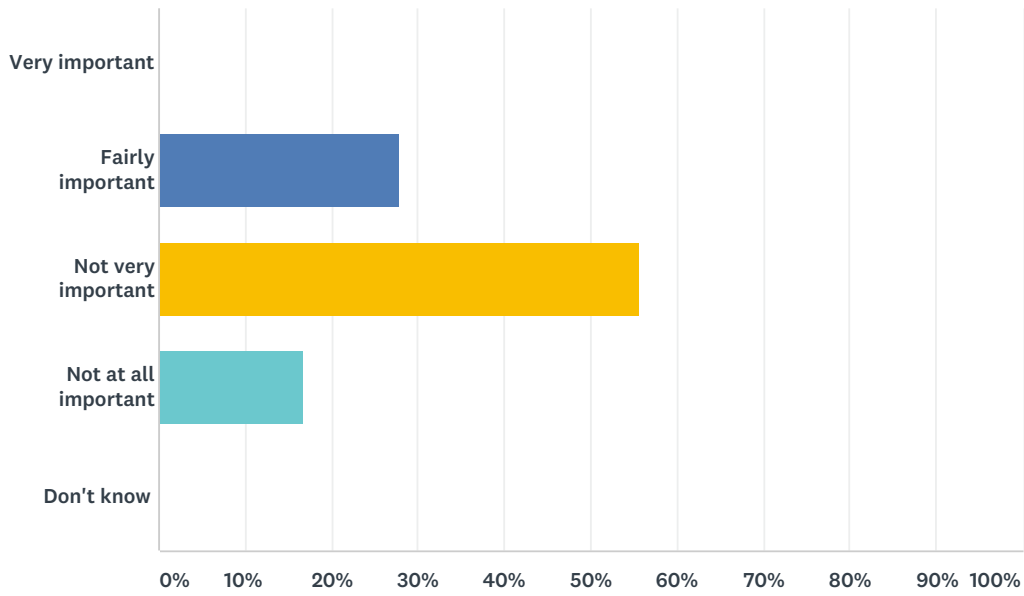
System A: 75% (below human expert level) accuracy, full explanation

System B: 85% (human expert level) accuracy, partial explanation

System C: 95% (beyond human level) accuracy, no explanation

Q1. How important is it for a kidney patient and their family to receive an explanation of an automated decision about why the patient could or could not be matched?

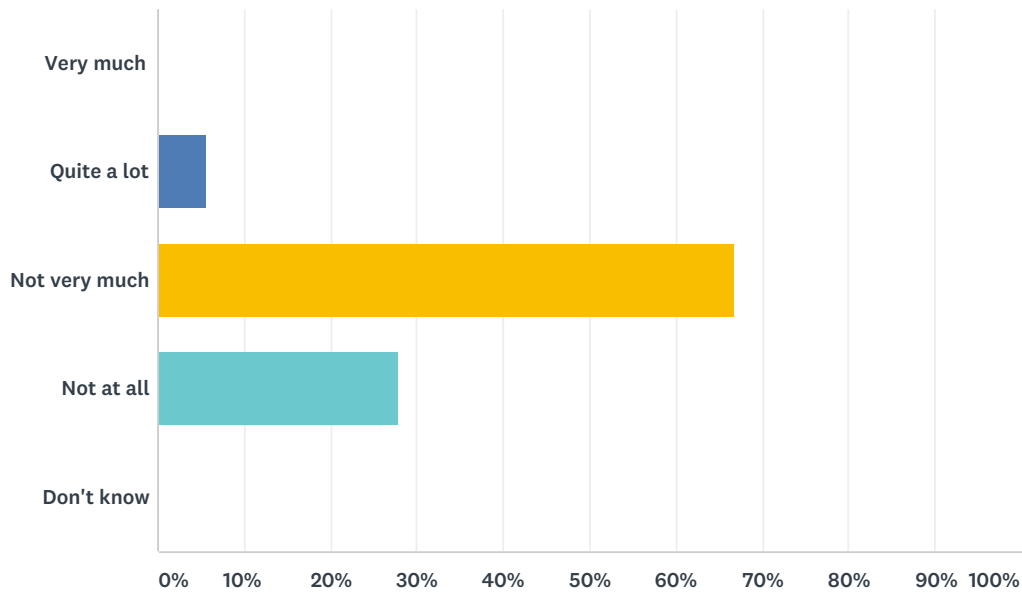
Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------------|-----------|-----------|
| Very important | 0.00% | 0 |
| Fairly important | 27.78% | 5 |
| Not very important | 55.56% | 10 |
| Not at all important | 16.67% | 3 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

Q4 Q2 If system C was chosen by the NHS, almost no explanation would be provided. How much does this matter?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|-----------|
| Very much | 0.00% | 0 |
| Quite a lot | 5.56% | 1 |
| Not very much | 66.67% | 12 |
| Not at all | 27.78% | 5 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

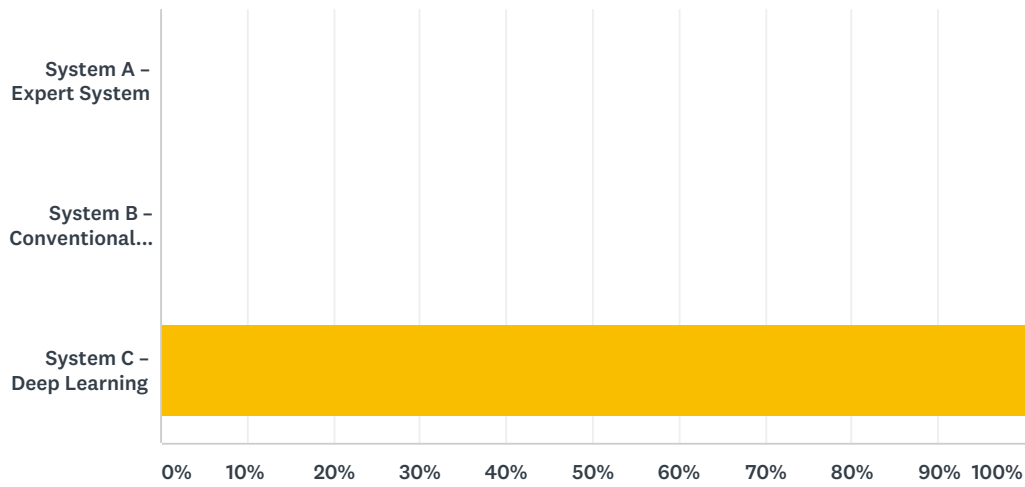
Q5 Why (up to three reasons)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Reason 1 | 100.00% | 18 |
| Reason 2 | 50.00% | 9 |
| Reason 3 | 38.89% | 7 |

Q6 Q3 Which automated decision system do you think the NHS should choose?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES |
|--|------------|
| System A – Expert System | 0.00% 0 |
| System B – Conventional machine learning | 0.00% 0 |
| System C – Deep Learning | 100.00% 18 |
| TOTAL | 18 |

Q7 Explain the factors affecting your choice (up to 3)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Factor 1 | 100.00% | 18 |
| Factor 2 | 61.11% | 11 |
| Factor 3 | 44.44% | 8 |

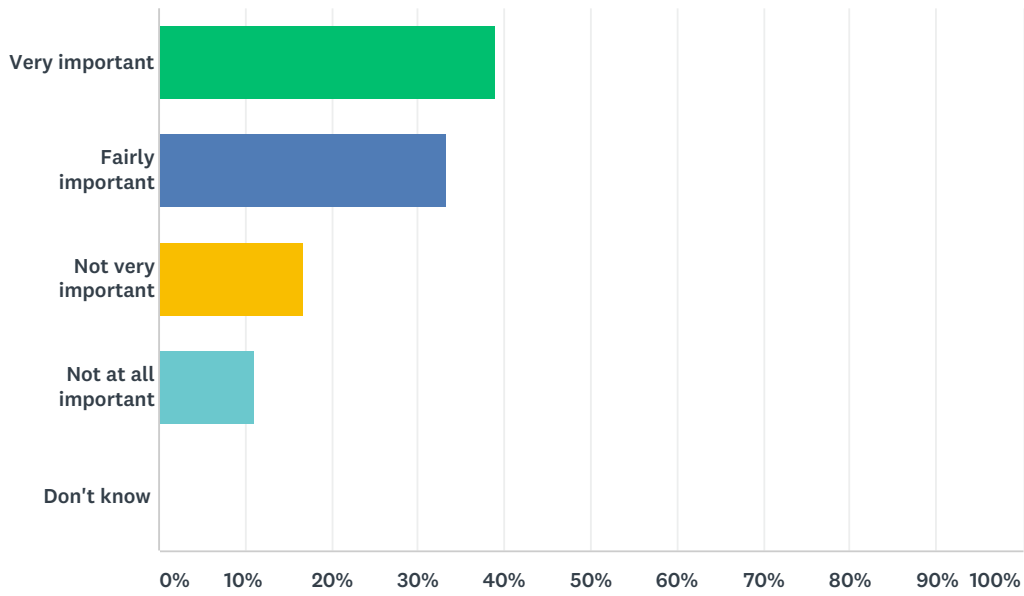
Q3 There are 3 automated decision systems for the Police to choose from:

- System A – Expert System This system uses an algorithm that was developed with help from very experienced Police Custody Officers, and aims to follow the same reasoning as they would do. In practice it does not reach the same level of accuracy as they would, but the algorithm is completely transparent in the way it reaches its conclusions: for each individual case it can provide specific rules that were applied to reach a conclusion. When tested on existing data about reoffending, this system was shown to have an overall accuracy rate of 75%. This means that 25% of the time its predictions were incorrect (e.g. predicting that an individual would commit a serious offence when in reality they didn't, or vice versa). The accuracy of this system is comparable to that of an average Police Custody Officer.
- System B – Conventional machine learning This system uses an algorithm that was established through machine learning from a large set of criminal offence data, collected by the police and local agencies. This algorithm achieves (human) expert level performance, but it is not very transparent in the way it reaches its conclusions: it can only tell us which features, in general, are important and which are not. When tested on existing data about reoffending, this system was shown to have an overall accuracy rate of 85%. This means that 15% of the time its predictions were incorrect (e.g. predicting that an individual would commit a serious offence when in reality they didn't, or vice versa.) The accuracy of this system is comparable to that of a very experienced Police Custody Officer.
- System C – Deep Learning This system uses advanced AI, derived from the same set of data as System B. However it has “taught itself” from the data. This algorithm is not transparent in the way it reaches conclusions: it is unable to provide any explanation that is understandable by humans. However it has an overall accuracy rate of 95%, which is better than human experts perform. This means that 5% of the time its predictions were incorrect (e.g. predicting that an individual would commit a serious offence when in reality they didn't, or vice versa.)

System A: 75% (Custody officer level) accuracy, full explanation
 System B: 85% (human expert level) accuracy, partial explanation
 System C: 95% (beyond human level) accuracy, no explanation

Q1. How important is it for an individual to receive an explanation of an automated decision about referral to a rehabilitation programme?

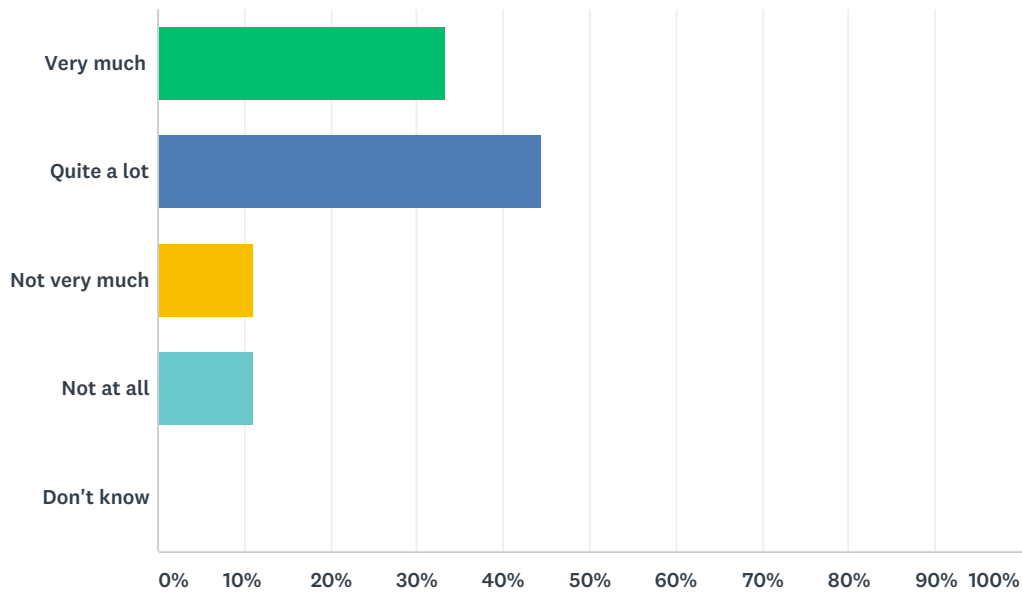
Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------------|-----------|-----------|
| Very important | 38.89% | 7 |
| Fairly important | 33.33% | 6 |
| Not very important | 16.67% | 3 |
| Not at all important | 11.11% | 2 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

Q4 Q2 If system C was chosen by the police force, almost no explanation would be provided. How much does this matter?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|-----------|
| Very much | 33.33% | 6 |
| Quite a lot | 44.44% | 8 |
| Not very much | 11.11% | 2 |
| Not at all | 11.11% | 2 |
| Don't know | 0.00% | 0 |
| TOTAL | | 18 |

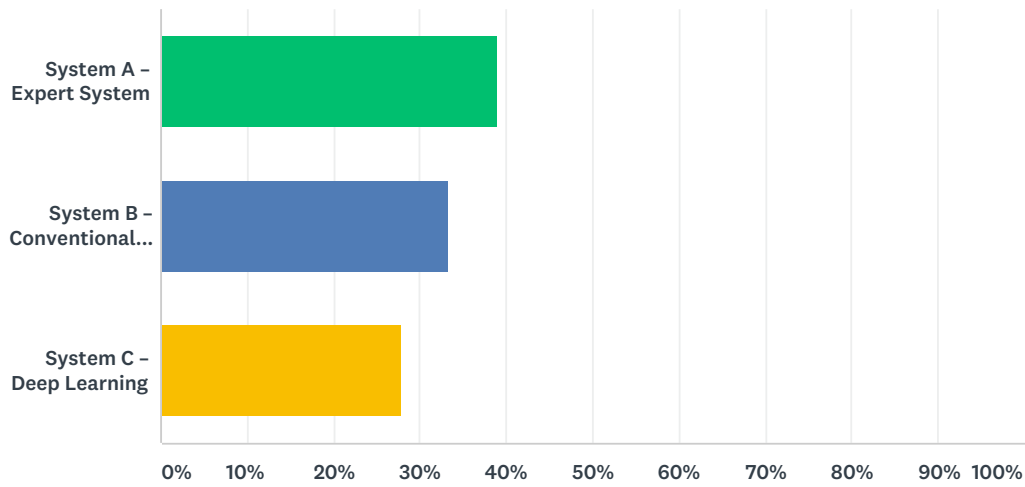
Q5 Why (up to three reasons)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Reason 1 | 100.00% | 18 |
| Reason 2 | 55.56% | 10 |
| Reason 3 | 38.89% | 7 |

Q6 Q3 Which automated decision system do you think the police force should choose?

Answered: 18 Skipped: 0



| ANSWER CHOICES | RESPONSES | |
|--|-----------|----|
| System A – Expert System | 38.89% | 7 |
| System B – Conventional machine learning | 33.33% | 6 |
| System C – Deep Learning | 27.78% | 5 |
| TOTAL | | 18 |

Q7 Explain the factors affecting your choice (up to 3)

Answered: 18 Skipped: 0

| ANSWER CHOICES | RESPONSES | |
|----------------|-----------|----|
| Factor 1 | 100.00% | 18 |
| Factor 2 | 72.22% | 13 |
| Factor 3 | 44.44% | 8 |